

Coloring random graphs online without creating monochromatic subgraphs¹

Torsten Mütze Thomas Rast

Institute of Theoretical Computer Science
ETH Zürich
8092 Zürich, Switzerland
{muetzet|trast}@inf.ethz.ch

Reto Spöhel[†]

Algorithms and Complexity Group
Max-Planck-Institut für Informatik
66123 Saarbrücken, Germany
rspoehel@mpi-inf.mpg.de

ABSTRACT. Consider the following random process: The vertices of a binomial random graph $G_{n,p}$ are revealed one by one, and at each step only the edges induced by the already revealed vertices are visible. Our goal is to assign to each vertex one from a fixed number r of available colors immediately and irrevocably without creating a monochromatic copy of some fixed graph F in the process.

Our first main result is that for any F and r , the threshold function for this problem is given by $p_0(F, r, n) = n^{-1/m_1^*(F, r)}$, where $m_1^*(F, r)$ denotes the so-called *online vertex-Ramsey density* of F and r . This parameter is defined via a purely deterministic two-player game, in which the random process is replaced by an adversary that is subject to certain restrictions inherited from the random setting. Our second main result states that for any F and r , the online vertex-Ramsey density $m_1^*(F, r)$ is a computable rational number.

Our lower bound proof is algorithmic, i.e., we obtain polynomial-time online algorithms that succeed in coloring $G_{n,p}$ as desired with probability $1 - o(1)$ for any $p(n) = o(n^{-1/m_1^*(F, r)})$.

CONTENTS

1. Introduction	2
1.1. The online setting	3
1.2. A general characterization of the online threshold	4
1.3. Remarks on Theorem 3	6
1.4. Remarks on Theorem 4	6
1.5. Algorithms for efficiently coloring random graphs online	7
1.6. Is there an explicit formula for $m_1^*(F, r)$?	7
1.7. Further related work	8
1.8. Organization of this paper	8

¹An extended abstract of this work has appeared in the proceedings of SODA '11.

[†]The author was supported by a fellowship of the Swiss National Science Foundation.

2. Proof ideas	9
2.1. Computing the online vertex-Ramsey density	9
2.2. About the proof of Theorem 3	12
2.3. About the proof of Theorem 4	13
3. Computing the online vertex-Ramsey density	13
3.1. Proof of Theorem 3	13
3.2. Definitions and notations	17
3.3. The algorithm	19
3.4. Simplifications and implementation of the algorithm	22
3.5. Basic properties of the algorithm	23
3.6. Graphs in $\mathcal{C}^{i,j,k}$ irrelevant for Builder	30
3.7. Further properties of the algorithm	34
4. Builder in the deterministic game	39
4.1. The pigeonholing	39
4.2. Builder's strategy and proof of Proposition 6	41
4.3. Analysis of ABSTRACTBUILD()	42
5. Painter in the deterministic game	47
5.1. Painter's strategy and proof of Proposition 7	47
5.2. A geometric viewpoint	49
5.3. Relation of the walk to other quantities	53
5.4. Analysis of PAINT()	57
5.5. Proof of Lemma 35	61
6. Proof of Theorem 4	65
6.1. Lower bound	65
6.2. Upper bound	67
Acknowledgement	69
References	70

1. INTRODUCTION

The study of colorability properties of random graphs has a rich history and has spurred many important developments in random graph theory. Thanks to the efforts of many researchers (e.g., [2, 3, 11, 14, 19, 29, 30, 35, 36]), very precise bounds on the chromatic number of the random graph are known by now. More recently, also several related coloring notions and their associated ‘chromatic numbers’ have been investigated for the random graph (e.g., [4, 9, 12, 23, 25, 28, 31, 40]).

In this work we are concerned with the following generalized notion of graph coloring: a coloring of a graph G is *valid* with respect to some given graph F if it contains no monochromatic copy of F , i.e., if there is no copy of F in G whose vertices all receive the same color. Note that a proper

coloring in the usual sense is a coloring that is valid with respect to a single edge. More generally, a coloring that is valid with respect to the star with ℓ rays is a coloring in which each color class induces a graph with maximum degree at most $\ell - 1$ (this is sometimes called an $(\ell - 1)$ -improper coloring, see [23] and references therein).

The main motivation for studying this notion of colorability comes from Ramsey theory, where one usually considers similarly defined *edge*-colorings. The threshold for the existence of a valid vertex-coloring of the random graph with respect to some given fixed graph F was determined by Łuczak, Ruciński, and Voigt [31]. To state their result, we introduce some terminology.

As usual, we denote by $G_{n,p}$ the random graph on n vertices (labelled from $1, \dots, n$) obtained by including each of the $\binom{n}{2}$ possible edges with probability $p = p(n)$ independently. We say that a graph G is (F, r) -*vertex-Ramsey* if every r -coloring of the vertices of G contains a monochromatic copy of F , i.e., if G does *not* allow a valid r -coloring with respect to F . A graph is a *matching* if its maximum degree is 1. We denote the number of edges and vertices of a graph H by $e(H)$ and $v(H)$, respectively.

Theorem 1 ([31]). *Let $r \geq 2$ be a fixed integer and F a fixed graph with at least one edge that in the case $r = 2$ is not a matching. Then there exist positive constants $c = c(F, r)$ and $C = C(F, r)$ such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}[G_{n,p} \text{ is } (F, r)\text{-vertex-Ramsey}] = \begin{cases} 0 & \text{if } p(n) \leq cn^{-1/m_1(F)} \\ 1 & \text{if } p(n) \geq Cn^{-1/m_1(F)} \end{cases},$$

where

$$m_1(F) := \max_{H \subseteq F: v(H) \geq 2} \frac{e(H)}{v(H) - 1}. \quad (1)$$

Note that the parameter $m_1(F)$ does not depend on r . It is widely believed (see e.g. [16]) that the threshold behaviour is even sharper than stated in Theorem 1. Friedgut and Krivelevich [18] proved this conjecture for the class of strictly 1-balanced graphs, i.e. for graphs F for which $e(H)/(v(H) - 1) < e(F)/(v(F) - 1)$ for all proper subgraphs $H \subsetneq F$ with $v(H) \geq 2$.

Implicit in the lower bound proof of Theorem 1 is the existence of a polynomial-time algorithm that a.a.s. succeeds in finding a valid coloring of $G_{n,p}$ for $p(n) \leq cn^{-1/m_1(F)}$, where polynomial here and throughout means polynomial in n for F and r fixed. (Here and throughout, a.a.s. stands for asymptotically almost surely, i.e., with probability tending to 1 as n tends to infinity.)

In this work, we study the same coloring problem in an *online setting*, and derive results of the same generality as those stated in Theorem 1 for the offline case.

1.1. The online setting. We consider the following online problem: The vertices of an initially hidden instance of $G_{n,p}$ are revealed one by one in increasing order, and at each step of the process only the edges induced by the vertices revealed so far are visible. Alternatively, one can think of the random edges leading from each vertex to previous vertices as being generated at the moment the vertex is revealed (each edge being inserted with probability p independently from all other edges). Each vertex has to be colored immediately and irrevocably with one of r available colors as soon as it is revealed, with the goal of avoiding monochromatic copies of a fixed graph F as before.

It follows from standard arguments (see [33, Lemma 7]) that this online problem has a threshold $p_0(F, r, n)$ in the following sense: For any function $p(n) = o(p_0)$ there is a strategy that a.a.s. finds an r -coloring of $G_{n,p}$ that is valid with respect to F online, and for any function $p(n) = \omega(p_0)$ any online strategy will a.a.s. fail to do so. (Observe that no computational restrictions are imposed in this definition, i.e., the coloring strategy is *not* required to be an efficient algorithm.)

Note that this is a weaker threshold behaviour than the one stated in Theorem 1. A closer inspection of the arguments in this paper shows that the online thresholds are indeed coarser than the offline

thresholds given by Theorem 1: the limiting probability that a valid coloring can be found online is a constant bounded away from 0 and 1 whenever $p(n)$ has the same order of magnitude as the threshold $p_0(F, r, n)$. This is a consequence of the fact that the online thresholds turn out to be determined by *local* substructures (see [22, Theorem 3.9]).

The online problem was first studied in [32], where the following simple strategy was analyzed. Assuming that the colors are numbered from 1 to r , the *greedy strategy* fixes an appropriate choice of subgraphs $H_1, \dots, H_r \subseteq F$, and at each step uses the highest-numbered color i that does not complete a monochromatic copy of H_i (or color 1 if no such color exists). Note that this strategy can easily be implemented in polynomial time.

For any graph F and any integer $r \geq 1$ we define the parameter $\overline{m}_1(F, r)$ recursively by

$$\overline{m}_1(F, r) := \begin{cases} \max_{H \subseteq F} \frac{e(H)}{v(H)} & \text{if } r = 1, \\ \max_{H \subseteq F} \frac{e(H) + \overline{m}_1(F, r-1)}{v(H)} & \text{if } r \geq 2. \end{cases} \quad (2)$$

The results of [32] can be stated as follows.

Theorem 2 ([32]). *For any fixed graph F with at least one edge and any fixed integer $r \geq 2$, the threshold for finding an r -coloring of $G_{n,p}$ that is valid with respect to F online satisfies*

$$p_0(F, r, n) \geq n^{-1/\overline{m}_1(F, r)},$$

where $\overline{m}_1(F, r)$ is defined in (2). A polynomial-time algorithm that succeeds a.a.s. for any $p(n) = o(n^{-1/\overline{m}_1(F, r)})$ is given by the greedy strategy.

If F has an induced subgraph $F^\circ \subsetneq F$ on $v(F) - 1$ vertices satisfying

$$m_1(F^\circ) \leq \overline{m}_1(F, 2), \quad (3)$$

where $m_1(F^\circ)$ is defined in (1), the greedy strategy is best possible, i.e., the threshold is

$$p_0(F, r, n) = n^{-1/\overline{m}_1(F, r)}.$$

The second part of Theorem 2 applies in particular to the case where F is a clique or a cycle of arbitrary fixed size (in these specific cases, the appropriate choice of subgraphs $H_1, \dots, H_r \subseteq F$ for the greedy strategy is $H_1 = \dots = H_r = F$). Thus for cliques and cycles, explicit threshold functions are known.

It was also pointed out in [32] that the greedy strategy is *not* best possible in general — as we shall see below, the threshold is significantly higher than what is guaranteed by Theorem 2 already in the innocent-looking case when F is a long path.

Our main result gives a combinatorial characterization of the online threshold that allows us to compute, for any F and r , a value $\gamma = \gamma(F, r)$ such that the threshold is given by $p_0(F, r, n) = n^{-\gamma}$. We also obtain polynomial-time coloring algorithms that a.a.s. find valid colorings of $G_{n,p}$ online in the entire regime below the respective thresholds, i.e., for any $p(n) = o(p_0)$.

1.2. A general characterization of the online threshold. Our main result characterizes the general threshold for the online problem in terms of a *deterministic two-player game*, which we describe in the following. The two players are called *Builder* and *Painter*, and the board is a graph that grows in each step of the game. Painter wants to maintain a valid coloring of the board, and her opponent Builder tries to prevent her from doing so by forcing her to create a monochromatic copy of F .

The game starts with an empty board, i.e., no vertices are present at the beginning of the game. In each step, Builder presents a new vertex and a number of edges leading from previous vertices

to this new vertex. Painter has to color the new vertex immediately and irrevocably with one of r available colors, and as before she loses as soon as she creates a monochromatic copy of F . Note that so far this is the same setting as before, except that we replaced ‘randomness’ by the second player Builder. (Put differently, if Builder presents every possible edge with probability p independently, this is exactly the online process introduced above.) However, we additionally impose the restriction that Builder is not allowed to present an edge that would create a (not necessarily monochromatic) subgraph H with $e(H)/v(H) > d$ on the board, for some fixed real number d known to both players. In other words, Builder must adhere to the restriction that the evolving board B satisfies $m(B) \leq d$ at all times, where as usual we define

$$m(B) := \max_{H \subseteq B} \frac{e(H)}{v(H)} .$$

We will refer to this game as the *deterministic F -avoidance game with r colors and density restriction d* .

We say that *Builder has a winning strategy* in this game (for a fixed graph F , a fixed number of colors r , and a fixed density restriction d) if he can force Painter to create a monochromatic copy of F within a finite number of steps. Conversely, we say that *Painter has a winning strategy* if she can avoid creating a monochromatic copy of F for an arbitrary number of steps. Note that if for some fixed F and r , Builder has a winning strategy for some density restriction d , then he also has a winning strategy for every density restriction $d' \geq d$. We say that a Painter or Builder strategy is *optimal* if it is a winning strategy simultaneously for all d for which the respective player has a winning strategy.

For any graph F and any integer $r \geq 2$ we define the *online vertex-Ramsey density* $m_1^*(F, r)$ as

$$m_1^*(F, r) := \inf \left\{ d \in \mathbb{R} \mid \begin{array}{l} \text{Builder has a winning strategy in the deterministic} \\ \text{\textit{F}-avoidance game with } r \text{ colors and density restriction } d \end{array} \right\} . \quad (4)$$

It is not hard to see that $m_1^*(F, r)$ is indeed well-defined for any F and r . With these definitions in hand, our results can be stated as follows.

Theorem 3. *For any graph F with at least one edge and any integer $r \geq 2$, the online vertex-Ramsey density $m_1^*(F, r)$ is a computable rational number, and the infimum in (4) is attained as a minimum.*

Theorem 4. *For any fixed graph F with at least one edge and any fixed integer $r \geq 2$, the threshold for finding an r -coloring of $G_{n,p}$ that is valid with respect to F online is*

$$p_0(F, r, n) = n^{-1/m_1^*(F, r)} , \quad (5)$$

where $m_1^*(F, r)$ is defined in (4). A polynomial-time algorithm that succeeds a.a.s. for any $p(n) = o(p_0)$ can be derived from one of Painter’s optimal strategies in the deterministic two-player game.

Theorem 4 reduces the problem of determining the threshold of the original probabilistic problem to the purely deterministic combinatorial problem of computing $m_1^*(F, r)$ or, informally speaking, of ‘solving’ the deterministic two-player game. According to Theorem 3, the latter is possible by a finite computation; note that in the asymptotic setting of Theorem 4, this is in fact a constant-size computation.

It follows from the results of [32] that for any graph F we have

$$\lim_{r \rightarrow \infty} m_1^*(F, r) = m_1(F) .$$

Thus the online thresholds approach the offline threshold given by Theorem 1 as the number of colors increases. It also follows from an example presented in [32] that if F is the disjoint union of two graphs H_1 and H_2 , the parameter $m_1^*(F, r)$ may be strictly higher than both $m_1^*(H_1, r)$ and $m_1^*(H_2, r)$. (Such a behaviour cannot occur for the parameter $m_1(F)$ appearing in Theorem 1.)

1.3. Remarks on Theorem 3. To put Theorem 3 into perspective, we mention that none of its three statements (computable, rational, infimum attained as minimum) is known to hold for the offline counterpart of $m_1^*(F, r)$, i.e., for the *vertex-Ramsey density*

$$m_1^o(F, r) := \inf \{m(G) \mid G \text{ is } (F, r)\text{-vertex-Ramsey}\}$$

introduced in [26]. It is also not known whether such statements are true for two analogous parameters related to *edge-colorings* (see [8, 27]). In fact, even the value of $m_1^o(P_3, 2)$ is unknown — the authors of [26] offer 400,000 zloty (Polish currency in 1993) for its exact determination (here P_3 denotes the path on three vertices).

As is the case for many parameters in Ramsey theory, computing the online vertex-Ramsey density $m_1^*(F, r)$ becomes intractable already for moderately large graphs F . We have an implementation that computes $m_1^*(F, 2)$ for all graphs F with at most 7 vertices in under 10 minutes on an ordinary desktop computer. Using more computational power, we managed to determine $m_1^*(F, 2)$ exactly for all non-forests on at most 9 vertices. (Our implementation is rather inefficient for forests — we believe that for this special case an adapted program would be much faster, see the remarks in Section 1.6 below.) The program can be downloaded from the authors' websites [1]. There might be some room for improvement here, but it seems unrealistic to compute $m_1^*(F, 2)$ for, say, general graphs with 20 vertices with our approach in reasonable time. (Recall that the Ramsey number $R(5)$, i.e., the smallest integer ℓ such that every edge-coloring of the complete graph on ℓ vertices contains a monochromatic K_5 , is still unknown!)

1.4. Remarks on Theorem 4. The intuition behind Theorem 4 is the following: It is well-known [10] (see also [22, Section 3]) that for any fixed graph G , a.a.s. the random graph $G_{n,p}$ with $p(n) = \omega(n^{-1/m(G)})$ contains a copy of G . The textbook second moment method proof of this fact can be adapted to show that for $p(n) = \omega(n^{-1/d})$ and any fixed finite Builder strategy for the deterministic two-player game that respects the density restriction d , a.a.s. the random process will exactly reproduce the given Builder strategy somewhere on the board. Thus if Builder has a winning strategy for a graph F and some given density restriction d , then in the probabilistic process with $p(n) = \omega(n^{-1/d})$, any online algorithm will a.a.s. be forced to create a monochromatic copy of F somewhere on the board. Consequently, the threshold of the probabilistic problem satisfies $p_0(F, r, n) \leq n^{-1/d}$. This argument is completely generic in the sense that it does not require any assumptions on the structure of Builder's winning strategy. The underlying connection between the probabilistic and a deterministic variant of the same problem was first observed in [8] (for the edge-coloring version of the problem studied here), and subsequently applied in [5]. The main contribution of the present work is that for the vertex-coloring problem studied here, the best upper bound on the online threshold resulting from this approach is tight (recall (4) and (5)).

By the argument we just sketched, every winning strategy for Builder in the deterministic game translates to an upper bound on the threshold of the probabilistic problem. It seems to be much harder to prove an equally general statement translating *Painter's* winning strategies in the deterministic game to *lower* bounds on the threshold of the probabilistic problem. The reason for this is that the probabilistic process satisfies the density restriction imposed on Builder only *locally*: Even though the random graph $G_{n,p}$ with $p(n) = \Theta(n^{-1/d})$ a.a.s. contains no *constant-sized* graphs G with $m(G) > d$, the density of larger subgraphs is unbounded — in particular, the expected density of the entire random graph is $\binom{n}{2}p/n = \Theta(n^{1-1/d})$, which is unbounded for $d > 1$. Consequently, winning strategies for Painter in the deterministic game do not automatically give rise to successful coloring strategies for the probabilistic problem. In order to nevertheless establish the desired lower bound result we will need a quite detailed understanding of the *structure* of Painter's and Builder's optimal strategies in the deterministic game.

Theorem 4 establishes a general correspondence between the original probabilistic problem and the deterministic two-player game. We are not aware of any other results that establish a similar correspondence between probabilistic and deterministic variants of the same problem. In particular, such a correspondence does not hold for the offline version of the problem studied here: according to Theorem 1, the threshold for the existence of a valid r -coloring w.r.t. F in the probabilistic setting is determined by the parameter $m_1(F)$, and not by the parameter $m_1^o(F, r)$ coming from the corresponding deterministic problem (in general we have $m_1(F) \neq m_1^o(F, r)$).

1.5. Algorithms for efficiently coloring random graphs online. We now describe the structure of the coloring algorithms that arise from our approach, and their relation to Painter's optimal strategies in the deterministic game.

We use the concept of *ordered graphs*. An ordered graph is a graph with an associated ordering of its vertices, where this ordering is interpreted as the order in which these vertices appeared in the probabilistic process or the deterministic game. We will see that for any graph F and any integer r , there exists an optimal Painter strategy (i.e., a strategy that is a winning strategy for any density restriction $d < m_1^*(F, r)$) that can be represented as a *priority list over ordered monochromatic subgraphs of F* . Such a priority list is computed along with $m_1^*(F, r)$ in our approach, and encodes the relative 'level of danger' Painter associates with copies of a given ordered subgraph of F in a given color. In the asymptotic setting of the probabilistic process, determining such a priority list is a constant-size computation.

Given this priority list, Painter's strategy is the following: Whenever Builder presents a new vertex, Painter determines for each color the most dangerous ordered graph that would be completed if the new vertex were assigned this color, and then selects the color for which this most dangerous graph is least dangerous among all colors. (Observe that this requires Painter to memorize the order in which the vertices on the board arrived.) Note that this strategy based on a priority list can be easily implemented in polynomial time.

As we shall see, for any F and r we can compute a priority list such that the strategy represented by it is not only (i) a winning strategy for Painter in the deterministic game with density restriction d for any $d < m_1^*(F, r)$, but also (ii) a (polynomial-time) algorithm that succeeds a.a.s. in finding a valid coloring of $G_{n,p}$ online for any $p(n) = o(n^{-1/m_1^*(F, r)})$. (Recall from Section 1.4 that (i) does not automatically imply (ii)!)

1.6. Is there an explicit formula for $m_1^*(F, r)$? From Theorem 2 and Theorem 4 it follows that for any graph F that has an induced subgraph $F^\circ \subsetneq F$ on $v(F) - 1$ vertices satisfying (3), for any $r \geq 2$ the online vertex-Ramsey density $m_1^*(F, r)$ is given by $\overline{m}_1(F, r)$ as defined in (2). (Of course, this can also be proved directly by considering only the deterministic game.)

The question arises whether also for general graphs F the abstract definition of $m_1^*(F, r)$ in (4) can be replaced by an explicit formula, perhaps by suitably generalizing the definition (2). To address this question we point out some of the difficulties involved in the innocent-looking case where F is a long path. The results concerning this special case will be published separately in the companion paper [37].

We first present a simplified formulation of our results for the case where F is an arbitrary forest. Suppose d is of the form $d = (k - 1)/k$ for some integer $k \geq 2$. Then the restriction that Builder is not allowed to create a subgraph of density more than d is equivalent to requiring that Builder creates no cycles and no components (=trees) with more than k vertices. We call this game the *deterministic F -avoidance game with r colors and tree size restriction k* .

ℓ	2, \dots, 27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45
$k^*(P_\ell, 2)$	$2^2, \dots, 27^2$	791	841	902	961	1040	1089	1156	1225	1323	1376	1449	1521	1641	1699	1796	1856	1991	2057
$k^*(P_\ell, 2) - \ell^2$	0	7	0	2	0	16	0	0	0	27	7	5	0	41	18	32	7	55	32

TABLE 1. Exact values of $k^*(P_\ell, 2)$ for $\ell \leq 45$.

Corollary 5 (Forests). *For any fixed forest F with at least one edge and any fixed integer $r \geq 2$, the threshold for finding an r -coloring of $G_{n,p}$ that is valid w.r.t. F online is*

$$p_0(F, r, n) = n^{-1-1/(k^*(F, r)-1)},$$

where $k^*(F, r)$ is the smallest integer k such that Builder has a winning strategy in the deterministic F -avoidance game with r colors and tree size restriction k .

(Corollary 5 can also be proved directly by a much simpler proof than the general arguments in this work, reusing ideas of [8].)

It is not hard to see that $k^*(F, r)$ is indeed well-defined for any forest F and any integer $r \geq 2$. It follows from the results in [32] that for any tree F and any integer $r \geq 2$ the greedy strategy (with $H_1 = \dots = H_r = F$) is a winning strategy for Painter in the deterministic game with tree size restriction $k = v(F)^r - 1$, i.e., guarantees a lower bound of $k^*(F, r) \geq v(F)^r$.

For the rest of this section we focus on the case where $F = P_\ell$ is the path on ℓ vertices, and $r = 2$ colors are available. For this case our general procedure for computing $m_1^*(F, r)$ (or, equivalently if F is a forest, for computing $k^*(F, r)$) can be simplified considerably. We were able to compute $k^*(P_\ell, 2)$ for all $\ell \leq 45$. The resulting values are stated in Table 1, where the bottom row shows the difference $k^*(P_\ell, 2) - \ell^2$, i.e., by how much optimal Painter strategies can improve on the greedy lower bound $v(P_\ell)^2 = \ell^2$. The values in Table 1 and the corresponding optimal Painter strategies seem to follow no discernible pattern. In view of this, it does not seem very likely that there exists an explicit formula for $k^*(P_\ell, 2)$, let alone for the parameter $m_1^*(F, r)$ in general.

The values in Table 1 also raise the question by how much optimal strategies can improve on the greedy lower bound asymptotically as $\ell \rightarrow \infty$. We can show that $k^*(P_\ell, 2) = \Omega(\ell^{2.01})$, i.e., there exist Painter strategies that improve on the greedy lower bound by a factor polynomial in ℓ . On the other hand, we can prove an upper bound of $k^*(P_\ell, 2) = \mathcal{O}(\ell^{2.59})$, which shows that no superpolynomial improvement is possible [37].

1.7. Further related work. The question of finding valid *edge*-colorings of random graphs online was first considered by Friedgut *et al.*, who proved a threshold result for the case where F is a triangle and $r = 2$ colors are available [17]. In [33, 34], the greedy strategy was analyzed for the edge-coloring setting, and results similar to Theorem 2 were derived for the case of $r = 2$ colors. In [8], we presented the upper bound approach via deterministic two-player games discussed above; this approach was applied by Balogh and Butterfield to derive new upper bounds for the case where F is a triangle and $r = 3$ colors are available [5]. It would be very interesting to determine whether a general result analogous to Theorem 4 holds for the edge-coloring setting.

Various edge-coloring Builder-Painter games were studied in the context of deterministic Ramsey theory. The smallest number of moves Builder needs to win in the deterministic edge-coloring game without any restrictions is called the *online (size) Ramsey number of F* and was studied by many researchers [6, 7, 15, 21, 27, 38, 39]. Variants of the game where Builder is subject to various restrictions were studied in [13, 20, 24].

1.8. Organization of this paper. Before actually proving Theorem 3 and Theorem 4, we informally present the main ideas behind our proofs in Section 2. In Section 3 we describe our procedure

to compute the online vertex-Ramsey density $m_1^*(F, r)$ for any F and r . In this section we formulate two central propositions (Proposition 6 and Proposition 7 below) which together show that an optimal Builder strategy and an optimal Painter strategy for the deterministic game can be derived from this procedure. The proof of Theorem 3 is based on these two propositions and is also presented in Section 3.

In Section 4 we prove Proposition 6 by deriving an explicit Builder strategy from the procedure presented in Section 3, and in Section 5 we prove Proposition 7 by deriving an explicit Painter strategy from the same procedure. These two sections can be read independently from each other.

In Section 6 we finally turn to the original probabilistic problem and present the proof of Theorem 4. While the upper bound proof is completely self-contained and can be read independently from all other proofs, the lower bound proof relies very much on our analysis of the deterministic game in the preceding sections.

2. PROOF IDEAS

In this section we aim to give an informal description of the main ideas behind our proofs. We will first focus on our procedure for computing the online vertex-Ramsey density, and then briefly comment on the proofs of Theorem 3 and Theorem 4.

2.1. Computing the online vertex-Ramsey density. Throughout this section, we focus on the deterministic game and sketch the underlying ideas in our procedure for computing the online vertex-Ramsey density $m_1^*(F, r)$ for given F and r . Note that the following is *not* a proof sketch of Theorem 3 — rather, our goal in this section is to develop some intuition for how one arrives at the key definitions which stand at the very beginning of our formal arguments.

2.1.1. Basic observations. Consider a family $\{G_1, \dots, G_f\}$ of disjoint copies of the same graph G on the board, and suppose that Builder adds vertices v_1, \dots, v_f to the board connecting v_i to G_i in exactly the same way for all $1 \leq i \leq f$. Then, by the pigeonhole principle, for a $(1/r)$ -fraction of the new vertices, Painter's coloring decision will be the same and result in copies of the same r -colored graph G^+ . By performing this pigeonholing in each step of his strategy, Builder can thus force Painter to always create *many* copies of one of the r graphs G^+ that Painter may choose from. Consequently, in the following we may assume w.l.o.g. that whenever Builder manages to enforce an r -colored graph G^+ on the board, he has as many such copies available as he needs in further steps.

As it turns out, the only type of move that is useful for Builder is of the following form: Assume that for each of the colors $s \in [r]$ the board contains a monochromatic copy of some subgraph H_s of F in color s . Then Builder can force Painter to extend one of these copies to a monochromatic copy of a subgraph H_σ^+ of F with $v(H_\sigma^+) = v(H_\sigma) + 1$ for a color $\sigma \in [r]$ by presenting a new vertex v and connecting it appropriately to the already existing monochromatic copies of H_1, \dots, H_r (see Figure 1). Furthermore, w.l.o.g. Builder will always perform such a step using monochromatic copies of the graphs H_1, \dots, H_r that have evolved independently from each other so far, and that are therefore contained in distinct components of the board (playing like this throughout will not increase the density restriction d for which Builder's strategy is legal). Proceeding in this fashion, Builder step by step enforces larger monochromatic subgraphs of F from smaller ones, and eventually a monochromatic copy of F (if the density restriction allows it).

Each monochromatic copy of some subgraph H of F created in this way is contained in a larger 'history graph' G that encodes all of Builder's construction steps that lead to the monochromatic copy of H . Using the notation from the preceding paragraph, the history graph G of H_σ^+ arises as the union of the history graphs G_1, \dots, G_r of the copies of H_1, \dots, H_r (due to our assumption on

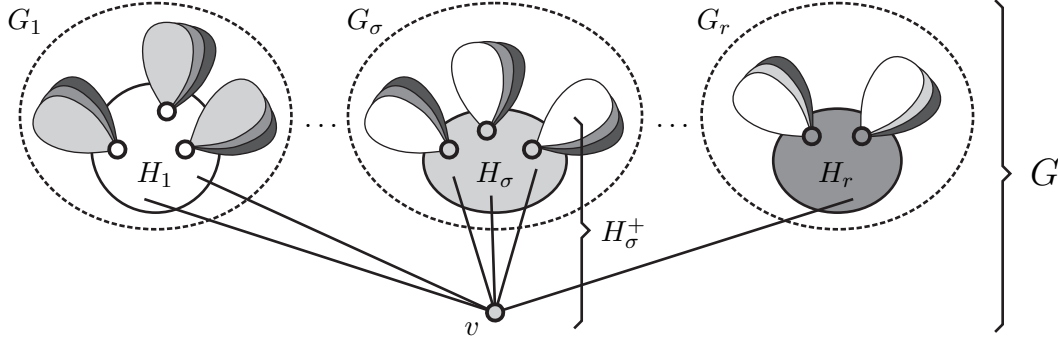


FIGURE 1. Builder enforces larger monochromatic subgraphs of F from smaller ones.

how Builder plays, these r history graphs are disjoint from each other), the vertex v and the edges that connect v to the copies of H_1, \dots, H_r in G_1, \dots, G_r .

2.1.2. Exploring Builder's options. The key ingredient in our approach is a systematic exploration from Builder's point of view which monochromatic subgraphs of F he can enforce against a *fixed* Painter strategy. Our final procedure for computing $m_1^*(F, r)$ will have to branch on different coloring decisions of Painter, each branching corresponding to a different Painter strategy, but these branchings do not interfere with the ideas we want to present here. In the following we therefore assume that Painter plays according to a fixed strategy, and explain on an intuitive level how Builder can determine the smallest density restriction d for which he can enforce a monochromatic copy of F against the given Painter strategy.

As a first approach to such a systematic exploration, Builder could maintain for each color $s \in [r]$ a list \mathcal{H}_s of all subgraphs H_s of F for which he has already enforced a monochromatic copy in color s against the given Painter strategy, and also record the specific way in which the graph H_s can be enforced by storing the corresponding history graph G_s . Builder can then use entries $(H_s, G_s) \in \mathcal{H}_s$, one for every color $s \in [r]$, to create new entries $(H_\sigma^+, G) \in \mathcal{H}_\sigma$ in the manner described above (see Figure 1), and compute for each such step the smallest density restriction d for which this step is legal. (Recall that by appropriate pigeonholing in each step, Builder can create as many copies of each entry as he needs on the board.) There is no obvious termination criterion for this procedure, i.e., without further arguments Builder can never be sure that he found the *smallest* possible density restriction d for which he can enforce a monochromatic copy of F against Painter's fixed strategy (it could be that by building larger and larger graphs he discovers new ways to enforce F that are compliant with smaller and smaller density restrictions). In the following we will sketch how this approach can be refined to eventually yield a procedure which is guaranteed to find the smallest such d in a finite number of steps.

2.1.3. A generalized density restriction. Note that each new history graph G arising in a given step of Builder has a *recursive* structure. Unfortunately, for computing the smallest admissible density restriction for which this step is legal the recursive structure of G does not help. However, by suitably generalizing our concept of density restriction the recursive structure of G can indeed be exploited.

For a fixed real number $\theta > 0$ and any graph H we define

$$\mu_\theta(G) := v(H) - e(H) \cdot \theta, \quad (6)$$

and consider the following generalization of the deterministic F -avoidance game with r colors and density restriction d : For fixed real parameters $\theta > 0$ and β we require that Builder adheres to the

restriction that every subgraph H of the evolving board B with $v(H) \geq 1$ satisfies

$$\mu_\theta(H) \geq \beta. \quad (7)$$

We refer to this game as the *deterministic F -avoidance game with r colors and generalized density restriction (θ, β)* . For any graph F with at least one edge, any integer $r \geq 2$ and any real number $\theta > 0$ we define the parameter

$$\beta^*(F, r, \theta) := \sup \left\{ \beta \in \mathbb{R} \left| \begin{array}{l} \text{Builder has a winning strategy in the deter-} \\ \text{ministic } F\text{-avoidance game with } r \text{ colors and} \\ \text{generalized density restriction } (\theta, \beta) \end{array} \right. \right\}. \quad (8)$$

Before discussing how this generalized game allows us to exploit the recursive structure of Builder's construction steps, let us explain how it relates to the original game with density restriction d that we are actually interested in.

Note that for any $\theta > 0$, the game with generalized density restriction $(\theta, 0)$ is equivalent to the game with density restriction $d = 1/\theta$. Together with the definition in (8) it follows that if for a given $\theta > 0$ we have $\beta^*(F, r, \theta) < 0$, then Painter has a winning strategy in the game with density restriction $d = 1/\theta$, and if $\beta^*(F, r, \theta) > 0$, then Builder has a winning strategy in the game with density restriction $d = 1/\theta$. So intuitively speaking, computing the online vertex-Ramsey density $m_1^*(F, r)$ is equivalent to determining the root of $\beta^*(F, r, \theta)$, although it is not clear yet whether such a root exists and whether it is unique.

As it turns out, $\beta^*(F, r, \theta)$ does indeed have a unique root $\theta^* = \theta^*(F, r)$, and the online vertex-Ramsey density $m_1^*(F, r)$ satisfies $m_1^*(F, r) = 1/\theta^*$. Furthermore, we can show that the root θ^* lies in an explicitly given finite set $Q = Q(F, r)$ of rational numbers. Therefore, it is straightforward to compute $m_1^*(F, r)$ provided we can compute $\beta^*(F, r, \theta)$ for given rational values of θ . We describe a procedure that does essentially that, with one major caveat: Observe that if $\beta \geq 0$, then the condition (7) holds for all subgraphs of the board if and only if it holds for all *connected* subgraphs. Our approach makes crucial use of this observation, and consequently our procedure computes $\beta^*(F, r, \theta)$ exactly for any input parameters F, r, θ for which $\beta^*(F, r, \theta) \geq 0$, but returns meaningless negative values on input parameters for which $\beta^*(F, r, \theta) < 0$. This makes no difference for our purposes since in order to find the root of $\beta^*(F, r, \theta)$ it suffices to check whether $\beta^*(F, r, \theta)$ equals zero for given values of $\theta \in Q$.

In the following we explain how the generalized density restriction allows us to exploit the recursive structure of the history graphs arising in the game. As before our viewpoint is that we are exploring Builder's options against a fixed strategy of Painter. More precisely, we consider a fixed value of $\theta > 0$, and our goal now is to determine the largest value of β for which Builder can enforce a monochromatic copy of F in the game with generalized density restriction (θ, β) against the given Painter strategy. Combining this with the already mentioned branching on different strategies of Painter allows us to compute $\beta^*(F, r, \theta)$ as defined in (8).

2.1.4. From history graphs to vertex weights. We return to considering Builder's construction step in which monochromatic copies of subgraphs H_1, \dots, H_r of F with the corresponding history graphs G_1, \dots, G_r are connected to a new vertex v , and Painter's decision to assign color σ to v creates a copy of H_σ^+ in color σ with history graph G (see Figure 1). In order to find the largest $\beta \geq 0$ for which this step is legal in the game with generalized density restriction (θ, β) , we need to find the minimal value $\mu_\theta(J)$ among all connected subgraphs J of G that contain v (recall that we may assume that J is connected due to the assumption that $\beta \geq 0$). As $\mu_\theta(J) = v(J) - e(J) \cdot \theta$ as defined in (6) is linear in $e(J)$ and $v(J)$, a connected subgraph J of G containing v that minimizes $\mu_\theta(J)$ can be found recursively as follows: determine *independently* for each $s \in [r]$ the connected subgraph J_s of $G_s + v$ containing v that minimizes $\mu_\theta(J_s)$, where $G_s + v$ denotes the subgraph of G induced by v and all vertices of the copy of G_s in G . The graph J we are interested in is then given by the union

of the graphs J_s for all $s \in [r]$. (Note that the subgraph J' of G that contains v and maximizes $e(J')/v(J')$ can *not* be found by independently considering each of the graphs $G_s + v$, $s \in [r]$ — this is precisely why we introduced the generalized notion of density restriction.) This independence allows us to compute the subgraph J of G that minimizes $\mu_\theta(J)$ *recursively* without remembering the actual structure of the history graphs G_s , $s \in [r]$. All the information that is necessary to do the same minimization in future steps (when the copy of H_σ^+ is extended to form larger subgraphs of F in color σ) can be stored by assigning the value $\sum_{s \in [r] \setminus \{\sigma\}} (\mu_\theta(J_s) - 1)$ to the vertex v in H_σ^+ (the -1 in the sum accounts for the fact that all the graphs $G_s + v$, $s \in [r]$, share the vertex v). In other words, we can condense the ‘history’ behind each of the vertices of a monochromatic copy of some subgraph of F into a single number. (Recall that we consider $\theta > 0$ to be fixed — this is crucial in all of the above.)

As a consequence, when maintaining the lists \mathcal{H}_s , $s \in [r]$, Builder no longer needs to store the entire history graph associated with some monochromatic subgraph H of F on one of these lists, but can store all the necessary information as a simple vertex-weighting of H . This greatly reduces the amount of information Builder needs to keep track of, but does not yet solve the issue that there is no obvious termination criterion for Builder’s exploration (Builder might still keep constructing new non-trivial entries forever).

2.1.5. Unique vertex weights via vertex orderings. In general, it may and will happen that the same subgraph H of F appears several times on one of Builder’s lists with different vertex-weightings in such a way that none of these entries is redundant — depending on how H is used in future steps, different vertex-weightings of the same graph might be desirable from Builder’s view. In other words, there is no unique best way of enforcing a copy of H in a given color for Builder.

It turns out, however, that different useful vertex-weightings can only arise if Builder presents the vertices of H in different orders (there are $v(H)!$ many different orders). For a fixed such order, there is a well-defined best vertex-weighting that Builder can achieve when enforcing H in that particular order. Thus to explore his options completely Builder only needs to compute *finite* lists \mathcal{H}_s , one for every color $s \in [r]$, which contain one entry for each vertex-ordering of every subgraph H of F .

This does not quite solve the issues we mentioned yet — it could still occur that Builder needs to recompute the vertex-weighting for a given entry many times because he finds better and better ways to enforce a given graph H in a particular order. To prevent this from happening, we need to be quite careful about the order in which we compute the entries of the lists \mathcal{H}_s — essentially we start by considering the game with generalized density restriction (θ, β) for the given fixed $\theta > 0$ and a very large β , and then successively lower β by the minimal amount that makes new options available to Builder. In each step we compute the weights for all graphs that Builder can create respecting the current generalized density restriction (θ, β) . This guarantees that we need to compute the weights for each graph only once, and therefore finally allows Builder to explore his options completely by a *finite* procedure.

2.1.6. Tying it all together. Along the lines sketched in the previous sections, we can compute $\beta^*(F, r, \theta)$ by dynamic programming over vertex-ordered subgraphs of F (provided that $\beta^*(F, r, \theta)$ is non-negative for the given $\theta > 0$, see the remarks in Section 2.1.3), branching on Painter’s decisions as appropriate. The online vertex-Ramsey density $m_1^*(F, r)$ can then be derived from $\beta^*(F, r, \theta)$ as explained in Section 2.1.3. As this is now a finite procedure, it also follows that the supremum in (8) is attained as a maximum, which with some further arguments also implies that the infimum in (4) is attained as a minimum.

2.2. About the proof of Theorem 3. For any graph F and any integer r , let $\tilde{m}(F, r)$ denote the value computed by the procedure outlined in Section 2.1. We prove that $\tilde{m}(F, r)$ equals $m_1^*(F, r)$

as defined in (4) by constructing explicit winning strategies for Builder and Painter, for arbitrary density restrictions $d \geq \tilde{m}(F, r)$ and $d < \tilde{m}(F, r)$, respectively.

For Builder such a strategy follows from the general principles underlying the procedure sketched in Section 2.1: all steps of the dynamic program which is at the heart of our approach can be interpreted as actual construction steps on the board of the deterministic game.

For Painter, such a strategy can be recovered from the branching on Painter's decisions performed in our procedure — we show that the decisions corresponding to a ‘worst’ path in the branching tree (viewed from Builder's perspective) give rise to a Painter strategy that succeeds in avoiding a monochromatic copy of F against *any* Builder strategy. This Painter strategy can be encoded by a priority list as described in Section 1.5.

To prove the success of this strategy, we use a *witness graph argument*: Essentially, we show inductively that whenever a monochromatic copy of some ordered subgraph H of F in some color $s \in [r]$ appears on the board, then this copy is contained in a graph that is at least as dense as indicated by the weights computed for H and the color s by the dynamic program in our procedure. (Recall from Section 2.1 that these weights basically encode the density of the history graph corresponding to the best way for Builder to enforce a monochromatic copy of H in color s .) This invariant holds in particular for all vertex-orderings of the graph F and all colors $s \in [r]$, and implies that whenever a monochromatic copy of F is completed, the board contains a graph that violates the density restriction imposed on Builder.

The proof of Theorem 3 we just sketched also shows that there exists an integer $a_{\max} = a_{\max}(F, r)$ such that for any given density restriction d Builder never needs more than a_{\max} steps to enforce a monochromatic copy of F , if he is able to do so at all. Note that this statement alone directly implies all three assertions of Theorem 3, as it shows that $m_1^*(F, r)$ can also be computed trivially by exhaustive search over the finitely many possible ways Builder and Painter can play in a_{\max} steps of the game.

2.3. About the proof of Theorem 4. We have already discussed the proof of the upper bound part of Theorem 4 in Section 1.4; as mentioned this proof is self-contained and does not depend on the rest of this work. The proof of the lower bound part is much more involved and relies on the same witness graph approach as the argument for Painter's success in the deterministic game described in the previous section. However, there is the additional issue that, as explained in Section 1.4, the random graph $G_{n,p}$ with $p(n) = o(n^{-1/m_1^*(F,r)})$ satisfies a density restriction of $d = m_1^*(F, r)$ only locally and *not* globally. Consequently, in order to apply the witness graph argument outlined above to the probabilistic setting of Theorem 4, we also need to show that the size of the witness graphs resulting from our arguments is bounded by some constant $v_{\max} = v_{\max}(F, r)$ (and not, say, linear in n). Unfortunately, we cannot show this for *all* priority lists that represent optimal strategies for Painter in the deterministic game. However, by applying a number of further technical refinements to the procedure described in Section 2.1, we can guarantee that it only computes priority lists for which a constant v_{\max} as desired indeed exists. It follows with the same witness graph argument as before that these priority lists represent polynomial-time coloring algorithms that a.a.s. succeed in finding a valid coloring of $G_{n,p}$ online for any $p(n) = o(n^{-1/m_1^*(F,r)})$.

3. COMPUTING THE ONLINE VERTEX-RAMSEY DENSITY

3.1. Proof of Theorem 3. Recall the definition of the deterministic F -avoidance game with r colors and generalized density restriction (θ, β) from Section 2.1.3, and recall further that, at least intuitively, computing the online vertex-Ramsey density $m_1^*(F, r)$ is equivalent to determining the root of $\beta^*(F, r, \theta)$ as defined in (8) (where existence and uniqueness of this root are not clear yet).

As already mentioned, we are going to derive a procedure that returns $\beta^*(F, r, \theta)$ for any $\theta > 0$ for which $\beta^*(F, r, \theta) \geq 0$, and a meaningless negative value for any $\theta > 0$ for which $\beta^*(F, r, \theta) < 0$. This procedure will be described in Section 3.3, and its output will be denoted by $\Lambda_\theta(F, r)$. We will see that the function $\Lambda_\theta(F, r)$ is well-defined for any real number $\theta > 0$, and for rational values of θ it can be computed using only integer arithmetic. Most of the remainder of this paper will be devoted to the proofs of the following two key statements.

Proposition 6 (Builder strategy from $\Lambda_\theta(F, r)$). *Let F be a graph with at least one edge and $r \geq 2$ an integer. There is a constant $a_{\max} = a_{\max}(F, r)$ such that the following holds: For any real numbers $\theta > 0$ and $\beta \geq 0$ with*

$$\Lambda_\theta(F, r) \geq \beta, \quad (9)$$

where $\Lambda_\theta()$ is defined in (24) below, Builder can enforce a monochromatic copy of F in the deterministic F -avoidance game with r colors and generalized density restriction (θ, β) in at most a_{\max} steps, regardless of how Painter plays.

Proposition 7 (Painter strategy from $\Lambda_\theta(F, r)$). *Let F be a graph with at least one edge, $r \geq 2$ an integer, and $\theta > 0$ and $\beta \geq 0$ real numbers such that*

$$\Lambda_\theta(F, r) < \beta, \quad (10)$$

where $\Lambda_\theta()$ is defined in (24) below.

Then Painter can avoid creating a monochromatic copy of F in the deterministic F -avoidance game with r colors and generalized density restriction (θ, β) , regardless of how Builder plays.

Before going into any details about the procedure that defines $\Lambda_\theta(F, r)$, we show how Proposition 6 and Proposition 7 imply Theorem 3.

For technical reasons, our formal arguments do not rely on the parameter $\beta^*()$ defined in (8), but on a related parameter that we introduce now. For any graph F with at least one edge, any integer $r \geq 2$, any real number $\theta > 0$ and any integer $a \geq a_{\min} := r(v(F) - 1) + 1$, we define

$$\beta'(F, r, \theta, a) := \sup \left\{ \beta \in \mathbb{R} \mid \begin{array}{l} \text{Builder has a winning strategy in the deterministic} \\ F\text{-avoidance game with } r \text{ colors and generalized} \\ \text{density restriction } (\theta, \beta) \text{ in at most } a \text{ steps} \end{array} \right\}. \quad (11)$$

Here the supremum is over a nonempty set of values because presenting the complete graph on a_{\min} vertices sequentially is a winning strategy for Builder that satisfies the generalized density restriction (θ, β) for any $\beta \leq \min\{k - \binom{k}{2} \cdot \theta \mid 1 \leq k \leq a_{\min}\}$. Note that for all F , r , and θ as before we have

$$\beta^*(F, r, \theta) = \sup_{a \geq a_{\min}} \beta'(F, r, \theta, a) = \lim_{a \rightarrow \infty} \beta'(F, r, \theta, a). \quad (12)$$

As in the definition of $\beta'()$ in (11) there is only a finite number of possible Builder strategies to consider, it is not hard to derive the following properties of $\beta'()$.

Lemma 8 (Properties of $\beta'(F, r, \theta, a)$). *For any graph F with at least one edge, any integer $r \geq 2$, any real number $\theta > 0$ and any integer $a \geq a_{\min}$, the supremum in (11) is attained as a maximum. For fixed F , r , and a as before, $\beta'(F, r, \theta, a)$ viewed as a function of $\theta > 0$ is continuous, non-increasing, piecewise linear, and has a unique root, which is contained in the set*

$$Q(a) := \{ 0 < \frac{v}{e} < 2 \mid v, e \in \mathbb{N} \wedge 1 \leq v \leq a \wedge 1 \leq e \leq \binom{v}{2} \}. \quad (13)$$

Proof. We identify Builder's strategies in the deterministic two-player game with r colors with finite r -ary rooted trees, where each node at depth k of such a tree is an r -colored graph on k vertices, representing the board after the k -th step of the game. Specifically, the tree \mathcal{T} representing a given Builder strategy is constructed as follows: The root of \mathcal{T} is the null graph (the graph whose vertex set is empty). The r children of any node B at depth k of \mathcal{T} are obtained by adding the $(k+1)$ -th

vertex of Builder's strategy to B (together with the edges that connect this vertex to previously added vertices according to Builder's strategy) and coloring it with one of the r available colors. Continuing like this, we construct \mathcal{T} , representing any situation in which Builder stops playing by a leaf of \mathcal{T} .

Note that in this formalization, a given tree \mathcal{T} represents a generic strategy for Builder (in the deterministic game with r colors) that may or may not satisfy a given generalized density restriction (θ, β) , and that can be thought of as a strategy for the ' F -avoidance' game for any given graph F . We say that \mathcal{T} is a *winning strategy* for Builder in a specific F -avoidance game if and only if every leaf of \mathcal{T} contains a monochromatic copy of F . We say that a Builder strategy \mathcal{T} is a *legal strategy* in the game with generalized density restriction (θ, β) if and only if (7) is satisfied for every subgraph H with $v(H) \geq 1$ of every node B in \mathcal{T} .

Let F , r and $a \geq a_{\min}$ be given. As the number of steps of the game is bounded by a , there is only a finite family $\mathfrak{T} = \mathfrak{T}(r, a)$ of different Builder strategies, obtained by exhaustive enumeration of all possible ways to add a new vertex to the board. Let $\mathfrak{W} = \mathfrak{W}(F, r, a) \subseteq \mathfrak{T}$ denote the set of winning strategies for Builder for the given F , and recall that for $a \geq a_{\min}$ the family \mathfrak{W} is nonempty.

Note that for any winning strategy $\mathcal{T} \in \mathfrak{W}$ and for any fixed $\theta > 0$,

$$f_{\mathcal{T}}(\theta) := \min_{\substack{B \in \mathcal{T} \\ H \subseteq B: v(H) \geq 1}} \mu_{\theta}(H) \quad (14)$$

is the maximal value of β such that \mathcal{T} is a legal strategy in the game with generalized density restriction (θ, β) . Optimizing over the (finite and nonempty) set of winning strategies, we obtain $\beta'(F, r, \theta, a)$ as defined in (11) as

$$\beta'(F, r, \theta, a) = \max_{\mathcal{T} \in \mathfrak{W}} f_{\mathcal{T}}(\theta) . \quad (15)$$

We conclude that the supremum in (11) is attained as a maximum. In the following we derive the claimed properties of $\beta'(F, r, \theta, a)$ as a function of $\theta > 0$ by considering the functions $f_{\mathcal{T}}(\theta)$, $\mathcal{T} \in \mathfrak{W}$.

Using (14) and combining the properties of the linear functions $\mu_{\theta}(H)$ for all $H \subseteq B$ with $v(H) \geq 1$ and all $B \in \mathcal{T}$ it is not hard to see that for any $\mathcal{T} \in \mathfrak{W}$ the function $f_{\mathcal{T}}(\theta)$ satisfies the following properties:

- $f_{\mathcal{T}}(\theta)$ is continuous and piecewise linear.
- There is an $\varepsilon = \varepsilon(\mathcal{T}) > 0$ such that $f_{\mathcal{T}}(\theta) = 1$ for all $0 < \theta \leq \varepsilon$ and $f_{\mathcal{T}}(\theta)$ is strictly decreasing for all $\theta \geq \varepsilon$.
- $f_{\mathcal{T}}(\theta)$ has a unique root in the set $\{\frac{v}{e} \mid v, e \in \mathbb{N} \wedge 1 \leq v \leq a \wedge 1 \leq e \leq \binom{v}{2}\}$.

Note that the root of $f_{\mathcal{T}}(\theta)$ is strictly smaller than 2: For any winning strategy $\mathcal{T} \in \mathfrak{W}$, there is a leaf B in \mathcal{T} that contains a (not necessarily monochromatic) copy of P_3 as a subgraph. This is trivially true if $P_3 \subseteq F$ (as every leaf of \mathcal{T} contains a monochromatic copy of F). If $P_3 \not\subseteq F$, then F is a matching, and any strategy where Painter colors endpoints of isolated edges on the board with different colors corresponds to a root-leaf path in \mathcal{T} that does not end with a matching (as otherwise \mathcal{T} would not be a winning strategy for Builder). Thus in either case the graph $H = P_3$ is a subgraph of some node B of \mathcal{T} , and consequently the minimization in (14) includes the function $\mu_{\theta}(H) = 3 - 2 \cdot \theta$, whose root is strictly smaller than 2.

It follows with (15) that also $\beta'(F, r, \theta, a)$ satisfies the three properties listed above, and that its root is strictly smaller than 2. Combining those properties shows that $\beta'(F, r, \theta, a)$ satisfies the conditions claimed in the lemma. \square

Using Proposition 6, Proposition 7 and Lemma 8, we will prove the following explicit version of Theorem 3.

Theorem 9 (Explicit Version of Theorem 3). *For any graph F with at least one edge and any integer $r \geq 2$, the online vertex-Ramsey density $m_1^*(F, r)$ defined in (4) satisfies*

$$m_1^*(F, r) = 1/\theta^* , \quad (16)$$

where $\theta^* = \theta^*(F, r)$ is the unique solution of

$$\Lambda_\theta(F, r) \stackrel{!}{=} 0 \quad (17)$$

and $\Lambda_\theta()$ is defined in (24) below.

Moreover, θ^* is a rational number from the set $Q(a_{\max})$, where $Q()$ is defined in (13) and a_{\max} is the constant guaranteed by Proposition 6. Furthermore, the infimum in (4) is attained as a minimum.

Theorem 3 is an immediate consequence of Theorem 9, observing that the solution of the equation (17) can be computed by evaluating $\Lambda_\theta(F, r)$ for all (finitely many) rational $\theta \in Q(a_{\max})$ (the constant a_{\max} is given explicitly in the proof of Proposition 6, see (97) below).

Proof of Theorem 9. Throughout the proof we consider F and r fixed and let $a_{\max} = a_{\max}(F, r)$ denote the constant guaranteed by Proposition 6.

Proposition 6 and Proposition 7 imply that for any given $\theta > 0$ for which $\Lambda_\theta(F, r) \geq 0$, the parameter $\Lambda_\theta(F, r)$ is the maximal value of β for which Builder can win the deterministic game with generalized density restriction (θ, β) , and if he can win then he needs at most a_{\max} steps to enforce a monochromatic copy of F , i.e., $\Lambda_\theta(F, r)$ coincides with $\beta'(F, r, \theta, a_{\max})$ as defined in (11).

Recall that according to Lemma 8 the supremum in (11) is always attained as a maximum, i.e., for any $\theta > 0$ Builder has a winning strategy in the game with generalized density restriction $(\theta, \beta'(F, r, \theta, a_{\max}))$. Thus if $\beta'(F, r, \theta, a_{\max}) \geq 0$ we must have that $\Lambda_\theta(F, r) \geq \beta'(F, r, \theta, a_{\max})$ as otherwise we could apply Proposition 7 with $\beta = \beta'$ to obtain a contradiction. Hence also $\Lambda_\theta(F, r)$ is non-negative in that case.

It follows that for any $\theta > 0$ the two functions $\Lambda_\theta(F, r)$ and $\beta'(F, r, \theta, a_{\max})$ either coincide or are both negative. Thus in particular they have the same set of roots, which by Lemma 8 consists of a single rational number $\theta^* = \theta^*(F, r)$ from the set $Q(a_{\max})$.

Applying Proposition 6 with $\theta = \theta^*$ and $\beta = 0$ yields that Builder has a winning strategy in the game with generalized density restriction $(\theta^*, 0)$ (in at most a_{\max} steps). Conversely, for any $\theta > \theta^*$ we obtain with Lemma 8 that $\beta'(F, r, \theta, a_{\max})$ is negative which, as discussed above, implies that also $\Lambda_\theta(F, r)$ is negative. Consequently we may apply Proposition 7 with $\beta = 0$ to infer that Painter has a winning strategy in the game with generalized density restriction $(\theta, 0)$.

Recalling that for any $\theta > 0$ the game with generalized density restriction $(\theta, 0)$ is equivalent to the original deterministic game with density restriction $d = 1/\theta$, we may restate our findings as follows: Builder has a winning strategy in the game with density restriction $d = 1/\theta^*$ (in at most a_{\max} steps), and for any $d < 1/\theta^*$ Painter has a winning strategy in the game with density restriction d . I.e., the online vertex-Ramsey density defined in (4) satisfies $m_1^*(F, r) = 1/\theta^*$, and the infimum in (4) is attained as a minimum. \square

Remark 10. Analogously to the second paragraph of the preceding proof it follows that for any given $\theta > 0$ for which $\Lambda_\theta(F, r) \geq 0$, also $\beta^*(F, r, \theta)$ as defined in (8) coincides with $\Lambda_\theta(F, r) = \beta'(F, r, \theta, a_{\max})$. Thus the unique root θ^* of $\Lambda_\theta(F, r) = \beta'(F, r, \theta, a_{\max})$ is also a root of $\beta^*(F, r, \theta)$.

Furthermore, the observation that the non-increasing functions $\beta'(F, r, \theta, a)$, $a \geq a_{\min}$, have a slope of at most -1 around their respective roots implies with (12) that the pointwise limit $\beta^*(F, r, \theta)$ has at most one root. Thus θ^* is indeed also the unique root of $\beta^*(F, r, \theta)$, as claimed in Section 2.1.3.

3.2. Definitions and notations. In order to present our procedure for computing the values $\Lambda_\theta(F, r)$ satisfying Proposition 6 and Proposition 7, we need to introduce a number of definitions and notations. Along with the definitions we give some intuition how those formal objects implement the ideas outlined in Section 2.1.

To simplify notation, for a graph H and any vertex v of H we abbreviate $v \in V(H)$ to $v \in H$. For a graph H and any set of vertices $U \subseteq V(H)$, we denote by $H \setminus U$ the graph obtained from H by removing all vertices in U and all edges incident to them. To indicate removal of a single vertex $v \in H$ we abbreviate $H \setminus \{v\}$ to $H \setminus v$.

3.2.1. Weighted graphs. A *vertex-weighted graph* is a graph H with a weight function $w : V(H) \rightarrow \mathbb{R}$. We refer to the values $w(u)$, $u \in H$, as *vertex weights*. Throughout this work, these vertex weights represent contributions to the linear function $\mu_\theta()$ defined in (6) that are obtained from ‘condensing’ history graphs as outlined in Section 2.1.4. They will always be non-positive.

For a fixed real number $\theta > 0$, any graph H , any vertex $v \in H$ and any weight function $w : V(H) \setminus \{v\} \rightarrow \mathbb{R}$ we define the value

$$d_\theta(H, v, w) := \min_{J \subseteq H : v \in J} \left(\sum_{u \in J \setminus v} (1 + w(u)) - e(J) \cdot \theta \right), \quad (18)$$

where the minimization is over all subgraphs J of H that contain the vertex v . As this minimization includes the graph J that consists only of the isolated vertex v , we always have $d_\theta(H, v, w) \leq 0$. Note that the minimum in (18) is always attained by an *induced* subgraph $J \subseteq H$. For convenience we will also use this notation for weight functions w whose domain is strictly larger than the set $V(H) \setminus \{v\}$. Of course, for the value of $d_\theta(H, v, w)$ only the values $w(u)$ of vertices $u \in H \setminus v$ are relevant.

The intuition behind the value $d_\theta(H, v, w)$ is the following: Assume that a copy of $H \setminus v$ is used as one of the graphs H_s in Figure 1, and Painter selects a color $\sigma \in [r]$ such that a copy of some other graph H_σ is extended to a copy of H_σ^+ . Then H becomes part of the history graph G of H_σ^+ , and the recursive contribution to the value $\mu_\theta(J)$ (as defined in (6)) of a subgraph $J \subseteq G$ minimizing $\mu_\theta(J)$ is exactly $d_\theta(H, v, w)$ if v is included in J . In our dynamic program, this will be recorded by adding a term of $d_\theta(H, v, w)$ to the vertex weight of v in H_σ^+ (and this is also how the vertex weights w of $H \setminus v$ were computed in earlier steps).

For a fixed real number $\theta > 0$, any graph H and any weight function $w : V(H) \rightarrow \mathbb{R} \cup \{-\infty\}$ we define

$$\lambda_\theta(H, w) := \sum_{u \in H} (1 + w(u)) - e(H) \cdot \theta. \quad (19)$$

As it is the case for the definition of $d_\theta()$ in (18), it is also convenient here to allow weight functions w whose domain is strictly larger than the set $V(H)$. Of course, for the value of $\lambda_\theta(H, w)$ only the values $w(u)$ of vertices $u \in H$ are relevant. Observe that $\lambda_\theta(H, w)$ defined in (19) can be written recursively for every vertex $v \in H$ as

$$\lambda_\theta(H, w) = \lambda_\theta(H \setminus v, w) + 1 + w(v) - \deg_H(v) \cdot \theta, \quad (20)$$

where $\deg_H(v)$ denotes the degree of v in H . This will be used several times in our arguments.

Using (6), $\lambda_\theta(H, w)$ defined in (19) can also be written as $\lambda_\theta(H, w) = \mu_\theta(H) + \sum_{u \in H} w(u)$, which intuitively means the following: If we imagine H to be at the center of a large history graph G , the parameter $\lambda_\theta(H, w)$ corresponds to the value $\mu_\theta(J)$ of the graph J obtained by attaching to each vertex $v \in H$ the $r - 1$ subgraphs that minimize $\mu_\theta(J_s)$ among all subgraphs J_s containing v in each of the $r - 1$ branches of the history graph G .

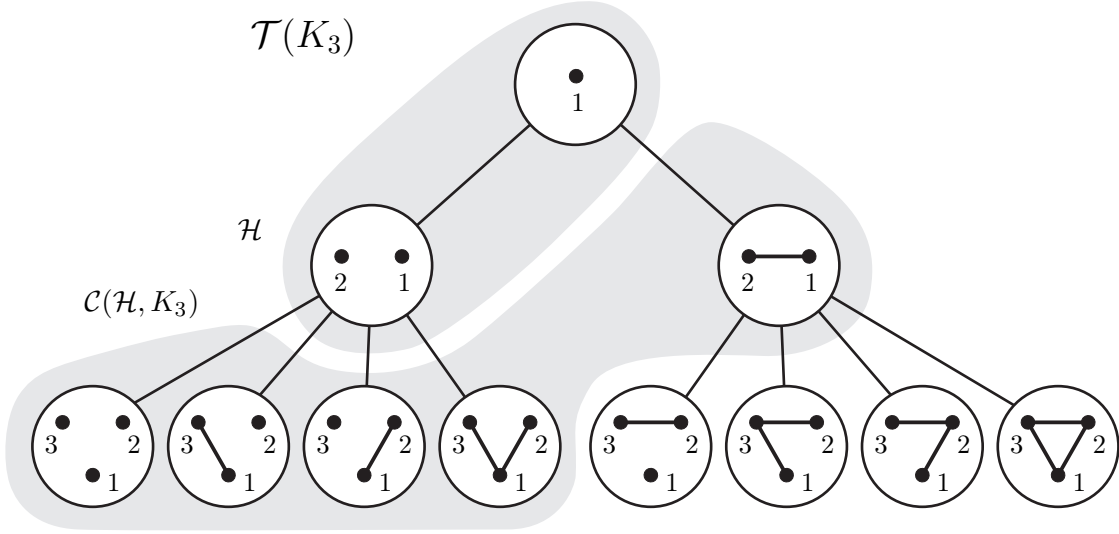


FIGURE 2. Illustration of the tree $\mathcal{T}(K_3)$ and the definition in (22) (the shaded regions represent subsets of nodes of $\mathcal{T}(K_3)$).

3.2.2. Ordered graphs. For any graph H , $h := v(H)$, a *vertex ordering* is a bijective mapping $\pi : V(H) \rightarrow \{1, \dots, h\}$, conveniently denoted by its preimages, $\pi = (\pi^{-1}(1), \dots, \pi^{-1}(h))$. An *ordered graph* is a pair (H, π) , where H is a graph and π is an ordering of its vertices. In the context of the F -avoidance game we interpret the ordering $\pi = (v_1, \dots, v_h)$ as the order in which the vertices of H appeared in the game, where v_h is the vertex that appeared first (we refer to it as the *oldest* vertex) and v_1 is the vertex that appeared last (we refer to it as the *youngest* vertex). We use $\Pi(V(H))$ to denote the set of all vertex orderings π of H .

For an ordered graph (H, π) and any subgraph $J \subseteq H$, we denote by $\pi|_J$ the order on the vertices of J induced by π . For any set $U \subseteq V(H)$ we use $\pi \setminus U$ as a shorthand notation for $\pi|_{H \setminus U}$. To indicate removal of a single vertex $v \in H$ we abbreviate $\pi \setminus \{v\}$ to $\pi \setminus v$.

Moreover, we define

$$\mathcal{S}(F) := \{[(H, \pi)]_{\sim} \mid H \subseteq F \text{ with } v(H) \geq 1 \text{ and } \pi \in \Pi(V(H))\} \quad (21)$$

as the family of all isomorphism classes of ordered subgraphs of F , where we write $(H, \pi) \sim (H', \pi')$ if (H, π) and (H', π') are isomorphic as ordered graphs. For simplicity we refer to the elements $[(H, \pi)]_{\sim} \in \mathcal{S}(F)$ in the following always as graphs $(H, \pi) \in \mathcal{S}(F)$. It is convenient to think of the graphs in $\mathcal{S}(F)$ as nodes of a rooted tree $\mathcal{T}(F)$ with root node $(K_1, (v_1))$ (an isolated vertex), where for each node $(H, \pi) \in \mathcal{S}(F)$, $\pi = (v_1, \dots, v_h)$, with $v(H) \geq 2$ the parent node is given by $(H \setminus v_1, \pi \setminus v_1)$. For any subset $\mathcal{H} \subseteq \mathcal{S}(F)$ we define the set $\mathcal{C}(\mathcal{H}, F) \subseteq \mathcal{S}(F)$ as

$$\mathcal{C}(\mathcal{H}, F) := \begin{cases} \{(K_1, (v_1))\} & \text{if } \mathcal{H} = \emptyset \\ \{(H, \pi = (v_1, \dots, v_h)) \in \mathcal{S}(F) \setminus \mathcal{H} \mid (H \setminus v_1, \pi \setminus v_1) \in \mathcal{H}\} & \text{otherwise} \end{cases} \quad (22)$$

Note that $\mathcal{C}(\mathcal{H}, F)$ is exactly the set of nodes of $\mathcal{T}(F)$ that are children of some node in \mathcal{H} , but that are not contained in \mathcal{H} . Figure 2 shows the tree $\mathcal{T}(F)$ for $F = K_3$ and illustrates the definition in (22).

Remark 11. Note that for two graphs $H_1 \subsetneq H_2$ with $v(H_1) = v(H_2)$, a monochromatic copy of H_1 on the board can never evolve into a copy of H_2 later in the game, as new edges appear only incident to newly added vertices. As a consequence, we could restrict our attention to *induced* subgraphs of F in all of our arguments. While changing the definition of $\mathcal{S}(F)$ in (21) accordingly would indeed lead

to some algorithmic savings (see Section 3.4), for our formal arguments we find it more convenient to include *all* subgraphs of F in the definition (21). Otherwise, unnecessary distraction would arise everytime an *induced* subgraph is mentioned in a proof.

3.3. The algorithm. In the following we present an algorithm `COMPUTEWEIGHTS()`, whose output is then used to define the function $\Lambda_\theta(F, r)$ that is referred to in Proposition 6 and Proposition 7.

Beside the graph F and the number of colors r , the algorithm has two more input parameters: the parameter θ from the generalized density restriction (see Section 2.1.3), and a finite sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ with the following interpretation: As indicated in Section 2.1, the underlying idea of the algorithm is to explore systematically from Builder's point of view which monochromatic ordered subgraphs of F he can enforce if Painter plays according to a fixed strategy. Step by step Builder enforces larger monochromatic subgraphs from smaller ones, and the appropriate vertex weights for these graphs are computed by dynamic programming. The sequence α encodes Painter's coloring decisions in the order they occur in the course of the algorithm (i.e., it represents a fixed Painter strategy), where an entry of this sequence may correspond to several coloring decisions of Painter for which she uses the same color. (Our proofs show that she would not gain anything by using different colors for these decisions.)

The algorithm maintains for each color $s \in [r]$ a family $\mathcal{H}_s \subseteq \mathcal{S}(F)$ of ordered subgraphs of F and a function $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$. The families \mathcal{H}_s correspond to the ordered subgraphs of F for which Builder has already enforced a monochromatic copy in color s . In the course of the algorithm, the families \mathcal{H}_s are successively enlarged. Initially, we have $\mathcal{H}_s = \emptyset$ for all $s \in [r]$, and at each step the candidate graphs to be added to the families \mathcal{H}_s are given by the sets $\mathcal{C}(\mathcal{H}_s, F)$ defined in (22); these correspond to the graphs that Builder can construct by adding a single vertex to a graph he has already enforced. Consequently, throughout the algorithm the families \mathcal{H}_s , $s \in [r]$, viewed as subsets of nodes of the tree $\mathcal{T}(F)$ defined after (21), grow downwards from the root.

For each $s \in [r]$ and each ordered graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, the function $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$ maintained by the algorithm induces a weight function $w_{(H, \pi, s)} : V(H) \rightarrow \mathbb{R}$ as follows: The weight $w_{(H, \pi, s)}(v_1)$ of the youngest vertex v_1 is given directly by $w_s(H, \pi)$; the weight $w_{(H, \pi, s)}(v_2)$ of the second-youngest vertex v_2 is given by $w_s(H \setminus v_1, \pi \setminus v_1)$, i.e., by the value of w_s for the parent of (H, π) in $\mathcal{T}(F)$; and so on. The full weight function $w_{(H, \pi, s)} : V(H) \rightarrow \mathbb{R}$ is therefore obtained by considering the value of w_s for all graphs on the path from (H, π) to the root $(K_1, (v_1))$ of the tree $\mathcal{T}(F)$, and each graph $(H, \pi) \in \mathcal{H}_s$ inherits all vertex weights except that of the youngest vertex from his ancestors in $\mathcal{T}(F)$.

More formally, and extending this construction to all graphs $(H, \pi) \in \mathcal{S}(F)$, we define for each $s \in [r]$ and each $(H, \pi) \in \mathcal{S}(F)$, $\pi = (v_1, \dots, v_h)$, the weight function

$$w_{(H, \pi, s)}(v_i) := \begin{cases} w_s(H \setminus \{v_1, \dots, v_{i-1}\}, \pi \setminus \{v_1, \dots, v_{i-1}\}) & \text{if } (H \setminus \{v_1, \dots, v_{i-1}\}, \pi \setminus \{v_1, \dots, v_{i-1}\}) \in \mathcal{H}_s, \\ -\infty & \text{otherwise.} \end{cases} \quad (23)$$

This notation will also be used in the formulation of `COMPUTEWEIGHTS()` in Algorithm 1 below. (We shall see that the algorithm never encounters the value $-\infty$ during its execution.) Note that an ordered graph $(H, \pi) \in \mathcal{S}(F)$ has vertices of weight $-\infty$ if and only if $(H, \pi) \in \mathcal{S}(F) \setminus \mathcal{H}_s$ for the corresponding $s \in [r]$, which intuitively means that Builder has not yet enforced a monochromatic copy of (H, π) in color s .

The families $\mathcal{H}_s \subseteq \mathcal{S}(F)$ and the functions $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$, $s \in [r]$, are extended step by step in the course of the algorithm, and their final values are returned when the algorithm terminates.

Algorithm 1: COMPUTEWEIGHTS(F, r, θ, α)

Input: a graph F with at least one edge, an integer $r \geq 2$, a real number $\theta > 0$, a sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$

Output: an r -tuple $((\mathcal{H}_s, w_s))_{s \in [r]}$, where $\mathcal{H}_s \subseteq \mathcal{S}(F)$ and $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$ for all $s \in [r]$

```

1 foreach  $s \in [r]$  do
2    $\mathcal{H}_s := \emptyset$ 
3    $\forall d \in \mathbb{R} : \mathcal{C}_s(d) := \emptyset$ 
4  $i := 0$ 
5 repeat (*)
6    $i := i + 1$ 
7   foreach  $s \in [r]$  do
8      $d_s^i := \max_{(H, \pi = (v_1, \dots, v_h)) \in \mathcal{C}(\mathcal{H}_s, F)} d_\theta(H, v_1, w_{(H, \pi, s)})$ 
9    $\sigma := \alpha_i$ 
10   $w^i := \sum_{s \in [r] \setminus \{s\}} d_s^i$ 
11   $j := 0$ 
12  repeat (**)
13     $j := j + 1$ 
14     $\mathcal{C}^{i,j} := \{(H, \pi = (v_1, \dots, v_h)) \in \mathcal{C}(\mathcal{H}_\sigma, F) \mid d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i\}$ 
15    if  $j = 1$  then
16       $\mathcal{C}_\sigma(d_\sigma^i) := \mathcal{C}^{i,1}$ 
17    foreach  $(H, \pi) \in \mathcal{C}^{i,j}$  do
18       $w_\sigma(H, \pi) := w^i$ 
19     $\mathcal{H}_\sigma := \mathcal{H}_\sigma \cup \mathcal{C}^{i,j}$ 
20     $k := 0$ 
21    repeat (***)
22       $k := k + 1$ 
23       $\mathcal{T}^{i,j,k} := \{(H, \pi = (v_1, \dots, v_h)) \in \mathcal{C}(\mathcal{H}_\sigma, F) \mid d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \geq d_\sigma^i\}$ 
24       $\mathcal{C}^{i,j,k} := \emptyset$ 
25      foreach  $(H, \pi) \in \mathcal{T}^{i,j,k}, \pi = (v_1, \dots, v_h)$ , do
26        if  $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) > d_\sigma^i$  or  $\nexists J \subseteq H : v_1 \in J \wedge (J, \pi|_J) \in \mathcal{C}_\sigma(d_\sigma^i)$  then
27          if  $\nexists J \subseteq H : v_1 \in J \wedge (J, \pi|_J) \in \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$  then
28             $\hat{i} := \max\{1 \leq \bar{i} \leq i \mid \alpha_{\bar{i}} = \sigma \wedge d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_{\sigma}^{\bar{i}}\}$ 
29          else
30             $\hat{i} := \max\{1 \leq \bar{i} \leq i \mid \alpha_{\bar{i}} = \sigma \wedge d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_{\sigma}^{\bar{i}}\}$ 
31             $w_\sigma(H, \pi) := w^{\hat{i}}$ 
32             $\mathcal{C}^{i,j,k} := \mathcal{C}^{i,j,k} \cup \{(H, \pi)\}$ 
33       $\mathcal{H}_\sigma := \mathcal{H}_\sigma \cup \mathcal{C}^{i,j,k}$ 
34    until  $\mathcal{C}^{i,j,k} = \emptyset$ 
35  until all  $(H, \pi) \in \mathcal{C}(\mathcal{H}_\sigma, F), \pi = (v_1, \dots, v_h)$ , satisfy  $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_\sigma^i$ 
36 until  $\mathcal{H}_s = \mathcal{S}(F)$  for some  $s \in [r]$ 
37 return  $((\mathcal{H}_s, w_s))_{s \in [r]}$ 

```

Consider now the algorithm `COMPUTEWEIGHTS()` as given in Algorithm 1. In the following we will try to convey an intuitive understanding of its operation, building on the informal remarks given in Section 2.1.

The algorithm works in rounds, and each round corresponds to relaxing the generalized density restriction (θ, β) by slightly lowering β , and then fully exploring Builder's options that become available as a consequence. Each iteration of the repeat-loop (*) is one such round.

At the beginning of the i -th round, for every color $s \in [r]$ the maximal $d_\theta()$ -value among all graphs in $\mathcal{C}(\mathcal{H}_s, F)$, denoted by d_s^i , is determined (lines 7–8). Here the sets $\mathcal{C}(\mathcal{H}_s, F)$ correspond to all graphs in color s that Builder could try to enforce next, and considering for each color a graph that *maximizes* $d_\theta()$ yields a new construction step for which β needs to be lowered *least* in order for that step to be compliant with the generalized density restriction (θ, β) . (Specifically, β needs to be lowered to $\beta_i := 1 + \sum_{s \in [r]} d_s^i$; note however that this successive lowering of β is not done explicitly in the algorithm.)

The i -th entry of the sequence α is then used to determine Painter's coloring decision $\sigma := \alpha_i$ for this construction step (line 9), and the rest of the round consists of updating the families \mathcal{H}_s and the functions w_s with all the information that can be extracted from that decision. In fact, only the family \mathcal{H}_σ grows; the families \mathcal{H}_s , $s \in [r] \setminus \{\sigma\}$, do not change.

The value $w^i := \sum_{s \in [r] \setminus \{\sigma\}} d_s^i$ defined in line 10 corresponds to the weight that needs to be assigned to the youngest vertex of every graph that is completed in color σ as a direct consequence of Painter's coloring decision. When the repeat-loop (**) is executed for the first time, those graphs are added to \mathcal{H}_σ via the set $\mathcal{C}^{i,1} \subseteq \mathcal{S}(F)$, and the function $w_\sigma : \mathcal{H}_\sigma \rightarrow \mathbb{R}$ is updated by assigning the value w^i to the newly created graphs (lines 14–19). For technical reasons, these graphs are also stored separately in a set $\mathcal{C}_\sigma(d_\sigma^i)$ that will be relevant later in the algorithm.

The remainder of the i -th round explores options that became available to Builder as a result of the graphs in $\mathcal{C}^{i,1}$ being added to \mathcal{H}_σ . These graphs can now be used themselves for further construction steps, and the graphs created in those construction steps can be used even further, etc. Some of these new potential construction steps are not legal for the current generalized density restriction (θ, β_i) , and will therefore only be explored in later rounds when (intuitively) β is lowered further. However, some of these are indeed legal for the current value of β , and it turns out that the previous decisions of Painter already imply which colors Painter should use in each of those construction steps (!). These indirect consequences of Painter's decision to use color $\sigma = \alpha_i$ in round i are explored in the repeat-loop (***), and in the repeat-loop (**) when it is executed for $j \geq 2$. The resulting graphs are added to \mathcal{H}_σ via the sets $\mathcal{T}^{i,j,k}, \mathcal{C}^{i,j,k} \subseteq \mathcal{S}(F)$ in the repeat-loop (***), and via the sets $\mathcal{C}^{i,j} \subseteq \mathcal{S}(F)$, $j \geq 2$, in the repeat-loop (**). This exploration of indirect consequences involves some technicalities for which we cannot give much intuition; see however the remarks in the first three paragraphs of Section 3.4 below. Note that the sets $\mathcal{C}_\sigma(d_\sigma^i)$ defined in line 16 (in this or an earlier round) come back into play in lines 26–27.

The i -th round terminates as soon as all ordered graphs $(H, \pi) \in \mathcal{C}(\mathcal{H}_\sigma, F)$, $\pi = (v_1, \dots, v_h)$, satisfy $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_\sigma^i$ (line 35). This corresponds to Builder having exhausted all his legal options in the game with generalized density restriction (θ, β_i) (recall that $\beta_i = 1 + \sum_{s \in [r]} d_s^i$). The $(i+1)$ -th round of the algorithm will then consider the game with generalized density restriction (θ, β_{i+1}) for some $\beta_{i+1} < \beta_i$.

The algorithm terminates as soon as one of the families \mathcal{H}_s , $s \in [r]$, contains all ordered subgraphs of F , i.e., $\mathcal{H}_s = \mathcal{S}(F)$ (line 36). This corresponds to Builder having enforced copies of all ordered subgraphs of F in color s (in particular, monochromatic copies of F in all possible vertex orderings).

We defer the formal arguments that `COMPUTEWEIGHTS()` is a well-defined algorithm and terminates correctly to Section 3.5, where we will prove the following claim.

Lemma 12 (Well-definedness and termination of algorithm). *All expressions that occur in the algorithm COMPUTEWEIGHTS() are well-defined, all numerical values and all sets that occur are finite, and on any input as specified the algorithm terminates correctly after at most $r \cdot |\mathcal{S}(F)|$ iterations of the repeat-loop (*).*

With Algorithm 1 in hand, we now define the parameter $\Lambda_\theta(F, r)$ for which we will prove Proposition 6 and Proposition 7.

For a fixed real number $\theta > 0$, any graph F with at least one edge and any integer $r \geq 2$ we define

$$\Lambda_\theta(F, r) := \min_{\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}} \max_{s \in [r]} \min_{\substack{H \subseteq F: v(H) \geq 1 \\ \pi \in \Pi(V(F))}} \lambda_\theta(H, w_{(H, \pi, s)}) , \quad (24)$$

where $\lambda_\theta()$ is defined in (19), and $w_{(H, \pi, s)}()$ is defined for all $(H, \pi) \in \mathcal{S}(F)$ and all $s \in [r]$ in (23) using the results $((\mathcal{H}_s, w_s))_{s \in [r]} := \text{COMPUTEWEIGHTS}(F, r, \theta, \alpha)$ of Algorithm 1.

We defer the formal arguments that $\Lambda_\theta(F, r)$ is well-defined to Section 3.5, where we will prove the following claim.

Lemma 13 (Well-definedness of $\Lambda_\theta(F, r)$). *For any real number $\theta > 0$, any graph F with at least one edge and any integer $r \geq 2$, the parameter $\Lambda_\theta(F, r)$ defined in (24) is a well-defined finite value.*

Note that for rational values of $\theta > 0$, the parameter $\Lambda_\theta(F, r)$ can be computed using only integer arithmetic.

Before we begin with the technical analysis of the algorithm COMPUTEWEIGHTS() in Sections 3.5—3.7, we give a few remarks about its implementation in the next section.

3.4. Simplifications and implementation of the algorithm. By Theorem 9, we can compute the online vertex-Ramsey density $m_1^*(F, r)$ as the inverse of the root of the parameter $\Lambda_\theta(F, r)$ defined in (24), where this definition involves the return values of the algorithm COMPUTEWEIGHTS(). As it turns out, the algorithm COMPUTEWEIGHTS() can be simplified considerably if one is only interested in computing the online vertex-Ramsey density $m_1^*(F, r)$ for given F and r (and not in proving Theorem 3 or Theorem 4, or in computing explicit winning strategies for Builder and Painter). The program to compute $m_1^*(F, r)$ that is available from the authors' websites [1] uses such a simplified version of the pseudocode above. In the following we outline the most important steps in this simplification.

First of all, the sets $\mathcal{C}_s(d)$ defined in line 16 and the case distinctions inside the repeat-loop (***) whether certain subgraphs are contained in those sets or not can be omitted, as they are only used for proving the lower bound part of Theorem 4, our result for the probabilistic problem. Specifically, these extra technicalities are needed to bound the size of the witness graphs for certain coloring strategies that are derived from the algorithm COMPUTEWEIGHTS() — recall from Section 2.2 and Section 2.3 that such a bound is unimportant for the deterministic game, but crucial for the original probabilistic problem (see also Lemma 35 and Remark 36 below).

In a second step the algorithm can be simplified even further: As it turns out, the entire repeat-loop (***) can be omitted; i.e., we do not need to compute any vertex weights for graphs $(H, \pi) \in \mathcal{C}(\mathcal{H}_\sigma, F)$, $\pi = (v_1, \dots, v_h)$, that satisfy $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) > d_\sigma^i$ for the current value of d_σ^i , and we do not need to add such graphs to the corresponding family \mathcal{H}_σ (thus for the color σ the algorithm will ignore the entire subtree of $\mathcal{T}(F)$ rooted at (H, π)). The reason for this is that such graphs are essentially useless for Builder, and therefore we do not need to consider them in our systematic exploration of Builder's options (see Lemma 23 and Algorithm 2 below).

Yet another simplification follows from Lemma 28 below: Combining (24) and (87) shows that we can change the return value of the algorithm COMPUTEWEIGHTS() to the sum on the right hand

side of (87) (note that this sum is exactly β_i , as used in our informal description of the algorithm). Thus we may stop the algorithm as soon as for some $\pi \in \Pi(V(F))$ the graph (F, π) is added to one of the families \mathcal{H}_s , $s \in [r]$, which may happen considerably earlier than the termination condition in line 36.

Further major savings are achieved by considering only induced subgraphs in the definition (21), as pointed out in Remark 11.

Some of these modifications might change the values returned by the algorithm `COMPUTEWEIGHTS`(F, r, θ, α) for a specific sequence α (as the families \mathcal{H}_s , $s \in [r]$, may evolve differently in the course of the algorithm, the entries of α get a different semantic), but not the value of $\Lambda_\theta(F, r)$ as defined in (24).

We conclude this section by sketching some ideas to further speed up the computation of $m_1^*(F, r)$ that do not directly relate to the pseudocode given in Algorithm 1.

When evaluating $\Lambda_\theta(F, r)$ for a given $\theta \in (0, 2)$, rather than calling the algorithm `COMPUTEWEIGHTS`() for each possible input sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ separately, we call it only once and, in each iteration, branch on all r values the variable σ can assume in line 9. Since most of the branches of the resulting recursion tree end after much fewer than $r \cdot |\mathcal{S}(F)|$ iterations, this allows us to evaluate the minimization in (24) and hence the value of $\Lambda_\theta(F, r)$ much more efficiently.

By Theorem 9 we have $m_1^*(F, r) = 1/\theta^*$, where $\theta^* = \theta^*(F, r)$ is the unique root of $\Lambda_\theta(F, r)$ defined in (24), which is guaranteed to be in the finite set $Q(a_{\max})$. In order to efficiently search for θ^* in $Q(a_{\max})$, we can exploit that the function $\Lambda_\theta(F, r)$ changes its sign from positive to negative at θ^* . Specifically, in order to compute θ^* , we alternate between shrinking the possible interval for the root θ^* by binary search (starting with the interval $(0, 2)$), and evaluating $\Lambda_\theta(F, r)$ for all rational values of θ inside the current interval up to a certain size of the denominator.

3.5. Basic properties of the algorithm. In this section we establish a number of basic properties of the algorithm `COMPUTEWEIGHTS`(), including several important monotonicity properties. We also provide the proofs for Lemma 12 and Lemma 13.

We begin by proving that the families \mathcal{H}_s grow downward from the root in the tree $\mathcal{T}(F)$ throughout the algorithm, as already mentioned.

Lemma 14 (Closure property of families \mathcal{H}_s). *Throughout the algorithm `COMPUTEWEIGHTS`() and for each $s \in [r]$ we have that if (H, π) , $\pi = (v_1, \dots, v_h)$, $h \geq 2$, is in \mathcal{H}_s , then $(H \setminus v_1, \pi \setminus v_1)$ is also in \mathcal{H}_s . In particular, if $\mathcal{H}_s \neq \emptyset$ then $(K_1, (v_1)) \in \mathcal{H}_s$.*

Proof. Observe that graphs are only added to \mathcal{H}_s in lines 19 and 33 of iterations for which $\alpha_i = s$, via the sets $\mathcal{C}^{i,j}$ and $\mathcal{C}^{i,j,k} \subseteq \mathcal{T}^{i,j,k}$. Thus by the definition of $\mathcal{C}^{i,j}$ in line 14 and of $\mathcal{T}^{i,j,k}$ in line 23, only graphs that are currently in $\mathcal{C}(\mathcal{H}_s, F)$ are added to \mathcal{H}_s . The claim thus follows from the definition of $\mathcal{C}(\mathcal{H}_s, F)$ in (22). \square

Next we prove the first part of Lemma 12, which states that all expressions that occur in the algorithm `COMPUTEWEIGHTS`() are well-defined and that all numerical values and all sets that occur are finite. (We ignore the assignment $\sigma := \alpha_i$ in line 9 for the time being — well-definedness of that assignment is immediate once we have proven the second part of Lemma 12, namely that the algorithm `COMPUTEWEIGHTS`() terminates correctly after at most $r \cdot |\mathcal{S}(F)|$ rounds.)

Proof of Lemma 12 (Well-definedness of algorithm). We need to argue that the expression $d_\theta(H, v_1, w_{(H, \pi, s)})$ (wherever it occurs) is well-defined, and that the maximum in line 8, line 28 and line 30 is always over a nonempty set.

Note that $w_\sigma(H, \pi)$ is defined in line 18 or line 31, just before (H, π) is added to \mathcal{H}_σ in line 19 or line 33, respectively. Combining this with Lemma 14 and using the definition in (23), it follows that the function $w_{(H, \pi, s)}$ defines finite vertex weights for all vertices of every graph $(H, \pi) \in \mathcal{H}_s$. Consequently, for every graph $(H, \pi) \in \mathcal{C}(\mathcal{H}_s, F)$, $\pi = (v_1, \dots, v_h)$, the function $w_{(H, \pi, s)}$ defines finite weights for *all but the youngest vertex* v_1 . As throughout the algorithm the function $d_\theta(H, v_1, w_{(H, \pi, s)})$ is evaluated only for graphs from the corresponding set $\mathcal{C}(\mathcal{H}_s, F)$, it follows that this expression (wherever it occurs) is indeed well-defined (recall the definition of $d_\theta()$ in (18)).

As long as none of the families \mathcal{H}_s , $s \in [r]$, contains all graphs from $\mathcal{S}(F)$, the sets $\mathcal{C}(\mathcal{H}_s, F)$ are nonempty. Thus the termination condition in line 36 ensures that the maximum in line 8 is always over a nonempty set and therefore well-defined.

For any $\sigma \in [r]$, let \tilde{i} be the smallest integer $i \geq 1$ for which $\alpha_i = \sigma$ and note that the \tilde{i} -th iteration of the repeat-loop (*) is the first one where graphs are added to the family \mathcal{H}_σ (initially, we have $\mathcal{H}_\sigma = \emptyset$). As $\mathcal{C}(\emptyset, F) = \{(K_1, (v_1))\}$ and $d_\theta(K_1, v_1, w_{(K_1, (v_1), \sigma)}) = 0$, by the definitions in lines 8, 14, and 16 we have $d_\sigma^{\tilde{i}} = 0$, and $(K_1, (v_1))$ is contained in $\mathcal{C}^{\tilde{i}, 1} = \mathcal{C}_\sigma(0)$. By the definition of $d_\theta()$ in (18), $d_\theta(H, v_1, w_{(H, \pi, \sigma)})$ on the right hand side of line 30 is always non-positive, i.e. less than or equal to $d_\sigma^{\tilde{i}} = 0$, implying that the maximum in line 30 is always over a nonempty set and therefore well-defined (this set contains at least the integer \tilde{i}). We can argue analogously for the set on the right hand side of line 28 if $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < 0$. On the other hand, if $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = 0$ then by what we said before the subgraph $J \subseteq H$ that consists only of the vertex v_1 is contained in $\mathcal{C}_\sigma(0) = \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$, and thus the condition in line 27 is violated.

For this last argument we used that no graphs are ever removed from the sets $\mathcal{C}_s(d)$, $d \in \mathbb{R}$. We have not shown this yet formally; however, it follows from Lemma 16 below that after the initialization in line 3, each set $\mathcal{C}_s(d)$ is modified at most once, namely in line 16 of the unique iteration i for which $\alpha_i = s$ and $d_s^i = d$. (If the reader is worried about this forward reference, he is welcome to substitute line 16 by $\mathcal{C}_\sigma(d_s^i) := \mathcal{C}_\sigma(d_s^i) \cup \mathcal{C}^{i, 1}$ for the time being, i.e., until Lemma 16 is proven.) \square

Having established that all numerical values assigned in the algorithm COMPUTEWEIGHTS() are finite, we state the following observations for further reference.

Lemma 15 (Finite and non-positive weights). *Throughout the algorithm COMPUTEWEIGHTS(), for each $s \in [r]$ we have that for each $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, all vertex weights $w_{(H, \pi, s)}(v_i)$ defined in (23) are finite non-positive values, while for each $(H, \pi) \in \mathcal{S}(F) \setminus \mathcal{H}_s$, at least one of the vertex weights is $-\infty$.*

Consequently, for all $(H, \pi) \in \mathcal{H}_s$, $\lambda_\theta(H, w_{(H, \pi, s)})$ defined in (19) is a finite value bounded by $v(F)$, while for all $(H, \pi) \in \mathcal{S}(F) \setminus \mathcal{H}_s$, we have $\lambda_\theta(H, w_{(H, \pi, s)}) = -\infty$.

Proof. By the definition of $d_\theta()$ in (18), $d_\theta(H, v_1, w_{(H, \pi, s)})$ on the right hand side of line 8 is always non-positive and finite, from which we conclude, using the definitions in line 10, 18 and 31, that $w_s(H, \pi)$ is a non-positive finite value for all $s \in [r]$ and all $(H, \pi) \in \mathcal{H}_s$. The first part of the statement now follows from the definition in (23) and Lemma 14. The second part follows from the first part using the definition in (19). \square

The next lemma establishes two important monotonicity properties, which will be used in many of the upcoming proofs.

Lemma 16 (Monotonicity of d_s^i and w^i in i). *Let $\sigma \in [r]$ and $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ be the input sequence of the algorithm COMPUTEWEIGHTS(). Throughout the algorithm, if $\alpha_i = \sigma$, then the variables d_s^i, d_s^{i+1} , $s \in [r]$, defined in line 8 satisfy*

$$d_\sigma^{i+1} < d_\sigma^i \quad \text{and} \quad d_s^{i+1} = d_s^i \quad \text{for all } s \in [r] \setminus \{\sigma\}.$$

Moreover, if $\alpha_{\bar{i}} = \alpha_i = \sigma$ for some $\bar{i} < i$, then the variables $w^{\bar{i}}, w^i$ defined in line 10 satisfy

$$w^i \leq w^{\bar{i}},$$

with equality if and only if $\alpha_{\bar{i}} = \alpha_{\bar{i}+1} = \dots = \alpha_i = \sigma$.

Of course, the variables d_s^i and d_s^{i+1} referred to in Lemma 16 are defined only if the number of iterations of the repeat-loop (*) is at least $i + 1$. Similarly, w^i and $w^{\bar{i}}$ are defined only if the number of iterations is at least i . Otherwise the statement of the lemma is void.

Proof. The first part of the lemma follows from the definition of d_σ^i in line 8, the termination condition in line 35, and the fact that none of the families \mathcal{H}_s , $s \in [r] \setminus \{\sigma\}$, is modified within the i -th iteration of the repeat-loop (*). The second part of the lemma follows from the first part and the definition of w^i in line 10. \square

For the proof that COMPUTEWEIGHTS() terminates correctly after at most $r \cdot |\mathcal{S}(F)|$ iterations of the repeat-loop (*) we will need the following auxiliary statement.

Lemma 17 (End of repeat-loop (**)). *Throughout the algorithm COMPUTEWEIGHTS(), at the end of each iteration of the repeat-loop (**), all graphs $(H, \pi) \in \mathcal{C}(\mathcal{H}_\sigma, F)$, $\pi = (v_1, \dots, v_h)$, satisfy*

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_\sigma^i. \quad (25)$$

For those graphs satisfying (25) with equality there is a subgraph $J \subseteq H$ with $v_1 \in J$ and $(J, \pi|_J) \in \mathcal{C}_\sigma(d_\sigma^i)$.

Proof. The j -th iteration of the repeat-loop (**) ends as soon as the repeat-loop (***) terminates. Due to the condition in line 34 this happens in the first iteration k for which the set $\mathcal{C}^{i,j,k} \subseteq \mathcal{T}^{i,j,k}$ is empty, which means that all graphs currently in $\mathcal{C}(\mathcal{H}_\sigma, F)$ violate the condition in the definition of $\mathcal{T}^{i,j,k}$ in line 23, or the condition in line 26. This implies the claim. \square

We now prove the second part of Lemma 12, namely that the algorithm COMPUTEWEIGHTS() terminates correctly after at most $r \cdot |\mathcal{S}(F)|$ rounds.

Proof of Lemma 12 (Termination of algorithm). Before bounding the number of iterations of the repeat-loop (*) we need to argue that the inner two repeat-loops always terminate.

Let $\sigma \in [r]$ and suppose that we have $\alpha_i = \sigma$ in the current iteration i of the repeat-loop (*). Note that in each iteration of the repeat-loop (***) except the last one, at least one element is added to the family \mathcal{H}_σ via the set $\mathcal{C}^{i,j,k}$ in line 33. Since throughout the algorithm, \mathcal{H}_σ is a subfamily of $\mathcal{S}(F)$, and since no graphs are ever deleted from \mathcal{H}_σ , the repeat-loop (***) terminates after at most $|\mathcal{S}(F)| + 1$ iterations.

It follows directly from the definition of d_σ^i in line 8 that in the first iteration $j = 1$ of the repeat-loop (**), the set $\mathcal{C}^{i,1}$ defined in line 14 is nonempty. As a consequence of the first part of Lemma 17 and the termination condition in line 35, the set $\mathcal{C}^{i,j}$ is also nonempty in all later iterations $j > 1$. Therefore, in each iteration of the repeat-loop (**), at least one element is added to the family \mathcal{H}_σ via the set $\mathcal{C}^{i,j}$ in line 19, and thus similarly to above the repeat-loop (**) terminates after at most $|\mathcal{S}(F)|$ iterations.

The above also implies that in each iteration of the repeat-loop (*), the size of exactly one of the families \mathcal{H}_s , $s \in [r]$, increases by at least one. Considering the condition in line 36, we conclude that the algorithm terminates after at most $r \cdot |\mathcal{S}(F)|$ iterations of the repeat-loop (*). \square

We proceed by proving Lemma 13, which states that $\Lambda_\theta()$ defined in (24) is a well-defined finite value.

Proof of Lemma 13. Due to the termination condition in line 36 of COMPUTEWEIGHTS(), for each possible input sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ there is an $s \in [r]$ for which the family \mathcal{H}_s returned by COMPUTEWEIGHTS() equals $\mathcal{S}(F)$. By Lemma 15, for this s the parameter $\lambda_\theta(H, w_{(H, \pi, s)})$ is a finite value for all $(H, \pi) \in \mathcal{S}(F)$. Thus for any fixed $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ the maximization over all $s \in [r]$ in (24) yields a finite value, and consequently also the outer minimization over all (finitely many) sequences $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ in (24) yields a finite value. \square

We continue with a simple invariant that holds at the beginning of each iteration of the main loop of COMPUTEWEIGHTS().

Lemma 18 (Beginning of repeat-loop (*)). *Throughout the algorithm COMPUTEWEIGHTS(), at the beginning of the i -th iteration of the repeat-loop (*), for every $s \in [r]$ all graphs $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, satisfy $d_\theta(H, v_1, w_{(H, \pi, s)}) > d_s^i$.*

Proof. Fix $i, s \in [r]$, and $(H, \pi) \in \mathcal{H}_s$ as in the lemma, and let $\bar{i} < i$ denote the iteration in which (H, π) was added to \mathcal{H}_s . We must have $\alpha_{\bar{i}} = s$, and (H, π) was added to \mathcal{H}_s either via one of the sets $\mathcal{C}^{\bar{i}, j}$ in line 19, or via one of the sets $\mathcal{C}^{\bar{i}, j, k} \subseteq \mathcal{T}^{\bar{i}, j, k}$ in line 33. By the conditions in the definition of $\mathcal{C}^{\bar{i}, j}$ in line 14 and of $\mathcal{T}^{\bar{i}, j, k}$ in line 23 we have $d_\theta(H, v_1, w_{(H, \pi, s)}) \geq d_s^{\bar{i}} > d_s^{\bar{i}+1} \geq d_s^i$, where the last two inequalities follow from the first part of Lemma 16. \square

The next lemma takes a closer look at the $d_\theta()$ -value of graphs that are added to the families \mathcal{H}_s via one of the sets $\mathcal{C}^{i, j, k}$ in the repeat-loop (**).

Lemma 19 (Sandwiched $d_\theta()$ -values for graphs in $\mathcal{C}^{i, j, k}$). *Let $\sigma \in [r]$ and $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ be the input sequence of the algorithm COMPUTEWEIGHTS(). Throughout the algorithm, if $\alpha_i = \sigma$, then for any graph (H, π) , $\pi = (v_1, \dots, v_h)$, that is added to the set $\mathcal{C}^{i, j, k}$ in line 32 (for this graph $w_\sigma(H, \pi)$ is defined in line 31 by setting it to $w^{\hat{i}}$) the following holds: If \hat{i} is defined in line 28, we have*

$$d_\sigma^{\hat{i}+1} \leq d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_\sigma^{\hat{i}}. \quad (26a)$$

Otherwise, i.e. if \hat{i} is defined in line 30, we have

$$d_\sigma^{\hat{i}+1} < d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_\sigma^{\hat{i}}. \quad (26b)$$

Proof. Clearly, the definition of \hat{i} in line 28 implies

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_\sigma^{\hat{i}}. \quad (27)$$

As (H, π) is in $\mathcal{C}^{i, j, k} \subseteq \mathcal{T}^{i, j, k}$, by the definition in line 23 we have $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \geq d_\sigma^i$, and consequently $\hat{i} < i$. Consider the smallest integer $\nu \geq 1$ for which $\alpha_{\hat{i}+\nu} = \alpha_{\hat{i}} = \sigma$, and note that $\hat{i} + \nu \leq i$. Again by the definition of \hat{i} in line 28 we have

$$d_\sigma^{\hat{i}+\nu} \leq d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_\sigma^{\hat{i}}. \quad (28)$$

Moreover, by the choice of ν and the first part of Lemma 16 we have

$$d_\sigma^{\hat{i}+\nu} = d_\sigma^{\hat{i}+1}. \quad (29)$$

Combining (28) and (29) yields the first inequality in (26a), and together with (27) proves the first part of the lemma. The second part follows similarly by using the definition of \hat{i} in line 30 and by interchanging $<$ with \leq in (27) and (28) (also in this case we must have $\hat{i} < i$, as otherwise we would have $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i$, a contradiction to the condition in line 26 and the negation of the condition in line 27). \square

The next lemma captures another important monotonicity condition: the lower the $d_\theta()$ -value of an ordered graph in a particular color is, the smaller is the weight assigned to its youngest vertex.

Lemma 20 ($d_\theta(\cdot)$ -value vs. weight monotonicity). *Throughout the algorithm COMPUTEWEIGHTS(), for every $s \in [r]$ and any two graphs $(H, \pi), (J, \tau) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, $\tau = (u_1, \dots, u_c)$, the following two properties hold:*

- If

$$d_\theta(H, v_1, w_{(H, \pi, s)}) < d_\theta(J, u_1, w_{(J, \tau, s)})$$

then we have

$$w_{(H, \pi, s)}(v_1) \leq w_{(J, \tau, s)}(u_1) .$$

- If

$$d_\theta(H, v_1, w_{(H, \pi, s)}) = d_\theta(J, u_1, w_{(J, \tau, s)}) \quad \text{and} \quad w_{(H, \pi, s)}(v_1) < w_{(J, \tau, s)}(u_1)$$

then $w_{(H, \pi, s)}(v_1) \stackrel{(23)}{=} w_s(H, \pi)$ is defined either in line 18 or in line 31 with \hat{i} defined in line 30, and $w_{(J, \tau, s)}(u_1) = w_s(J, \tau)$ is defined in line 31 with \hat{i} defined in line 28.

For the second part of Lemma 20 note that $w_s(H, \pi)$ and $w_s(J, \tau)$ are defined exactly once in the course of the algorithm (each in some iteration of the various repeat-loops), just before the corresponding graph (H, π) or (J, τ) is added to the family \mathcal{H}_s .

Proof. Let i_{\max} denote the total number of iterations of the repeat-loop (*). Fix some $\sigma \in [r]$ and consider the set $R_\sigma \in \mathbb{R}^2$ defined by

$$R_\sigma := \bigcup_{1 \leq i \leq i_{\max} : \alpha_i = \sigma} \left\{ ((1-t) \cdot d_\sigma^i + t \cdot d_\sigma^{i+1}, w^i) \mid t \in [0, 1] \right\} , \quad (30)$$

where we use the convention $d_\sigma^{i_{\max}+1} := d_\sigma^{i_{\max}}$ if $\alpha_{i_{\max}} = \sigma$. By Lemma 16 we have for any two pairs $(x, y), (x', y') \in R_\sigma$

$$x < x' \implies y \leq y' . \quad (31)$$

Now fix some graph $(H, \pi) \in \mathcal{H}_\sigma$, $\pi = (v_1, \dots, v_h)$, and consider the iteration i of the repeat-loop (*) where $\alpha_i = \sigma$ and where (H, π) is added to the family \mathcal{H}_σ . If (H, π) is added to \mathcal{H}_σ in line 19 then we have

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i \quad \text{and} \quad w_{(H, \pi, \sigma)}(v_1) \stackrel{(23)}{=} w_\sigma(H, \pi) = w^i \quad (32)$$

by the definitions in line 14 and line 18. If on the other hand (H, π) is added to \mathcal{H}_σ in line 33 then we obtain with Lemma 19 that

$$d_\sigma^{\hat{i}+1} \leq d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_\sigma^{\hat{i}} \quad \text{and} \quad w_{(H, \pi, \sigma)}(v_1) \stackrel{(23)}{=} w_\sigma(H, \pi) = w^{\hat{i}} \quad (33)$$

for some $\hat{i} \leq i$ with $\alpha_{\hat{i}} = \sigma$ (defined either in line 28 or in line 30).

Combining (30), (32) and (33) shows that any graph $(H, \pi) \in \mathcal{H}_\sigma$, $\pi = (v_1, \dots, v_h)$, satisfies

$$(d_\theta(H, v_1, w_{(H, \pi, \sigma)}), w_{(H, \pi, \sigma)}(v_1)) \in R_\sigma . \quad (34)$$

The first part of the claim now follows from (31) and (34).

The second part of the claim follows from (30), (34) and Lemma 16 by examining where in the set R_σ the point $(d_\theta(H, v_1, w_{(H, \pi, \sigma)}), w_{(H, \pi, \sigma)}(v_1))$ can possibly be located, depending on whether (H, π) is added to \mathcal{H}_σ in line 19 (then $w_\sigma(H, \pi)$ is defined in line 18) or in line 33 (then $w_\sigma(H, \pi)$ is defined in line 31), where the cases whether \hat{i} is defined in line 28 or in line 30 have to be distinguished using (26a) and (26b) from Lemma 19. \square

In the following we will repeatedly use the following auxiliary statement, which is an immediate consequence of the definition in (18).

Lemma 21 (Weights vs. $d_\theta(\cdot)$ -value monotonicity). *Let H be a graph, $v \in H$ and $J \subseteq H$ with $v \in J$. Moreover, let $w_H : V(H) \setminus \{v\} \rightarrow \mathbb{R}$ and $w_J : V(J) \setminus \{v\} \rightarrow \mathbb{R}$ with $w_H(u) \leq w_J(u)$ for all $u \in J \setminus v$. Then we have $d_\theta(H, v, w_H) \leq d_\theta(J, v, w_J)$.*

The next lemma establishes an important monotonicity condition for the vertex weights with respect to taking (ordered) subgraphs: the weights of a subgraph are always at least as high as the weights of the entire graph.

Lemma 22 (Subgraph weight monotonicity). *Throughout the algorithm COMPUTEWEIGHTS(), for every $s \in [r]$, if $(H, \pi) \in \mathcal{H}_s$, then for every subgraph $J \subseteq H$ we have $(J, \pi|_J) \in \mathcal{H}_s$ and $w_{(H, \pi, s)}(u) \leq w_{(J, \pi|_J, s)}(u)$ for all $u \in J$.*

Proof. We will prove the following auxiliary claim: for every $s \in [r]$, if $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, then for every subgraph $J \subseteq H$ with $v_1 \in J$ we have $(J, \pi|_J) \in \mathcal{H}_s$ and $w_{(H, \pi, s)}(u) \leq w_{(J, \pi|_J, s)}(u)$ for all $u \in J$. This implies the original claim, where the subgraphs $J \subseteq H$ are not required to contain the youngest vertex v_1 , as follows: if $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, and $J \subseteq H$ is *any* subgraph of H , then by Lemma 14 we also have $(H^{-c}, \pi^{-c}) \in \mathcal{H}_s$ where $c := \min\{i \mid v_i \in J\} - 1$ and $(H^{-c}, \pi^{-c}) := (H \setminus \{v_1, \dots, v_c\}, \pi \setminus \{v_1, \dots, v_c\})$. Moreover, $(J, \pi|_J)$ contains the youngest vertex v_{c+1} of (H^{-c}, π^{-c}) . Therefore, applying the auxiliary claim to (H^{-c}, π^{-c}) and $(J, \pi|_J)$, together with the observation that $w_{(H, \pi, s)}(u) = w_{(H^{-c}, \pi^{-c}, s)}(u)$ for all $u \in H^{-c}$ completes the argument.

To prove the auxiliary claim we argue by induction over the number of vertices of H . The claim clearly holds if H consists only of a single vertex, as then $J = H$ is the only subgraph of H . For the induction step let $\sigma \in [r]$ and consider a graph $(H, \pi) \in \mathcal{H}_\sigma$, $\pi = (v_1, \dots, v_h)$, with at least two vertices. We consider the iteration i of the repeat-loop (*) where $\alpha_i = \sigma$ and where (H, π) is added to the family \mathcal{H}_σ . Let J be a subgraph of H with $v_1 \in J$. By Lemma 14 we have that $(H \setminus v_1, \pi \setminus v_1) \in \mathcal{H}_\sigma$ at this point, and thus we know by induction that

$$(J \setminus v_1, \pi|_{J \setminus v_1}) \in \mathcal{H}_\sigma \quad (35)$$

and that

$$w_{(H, \pi, \sigma)}(u) \leq w_{(J, \pi|_J, \sigma)}(u) \quad \text{for all } u \in J \setminus v_1. \quad (36)$$

To complete the proof we only need to show two things: Firstly, that $(J, \pi|_J)$ is either already contained in \mathcal{H}_σ or added to this set together with (H, π) at the latest, and secondly, that $w_{(H, \pi, \sigma)}(v_1) \leq w_{(J, \pi|_J, \sigma)}(v_1)$ for the last vertex v_1 .

Recall that graphs are only added to \mathcal{H}_σ via one of the sets $\mathcal{C}^{i,j}$ in line 19, or via one of the sets $\mathcal{C}^{i,j,k} \subseteq \mathcal{T}^{i,j,k}$ in line 33. If (H, π) is contained in one of the sets $\mathcal{C}^{i,j}$, then by the definition in line 14 we have

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i, \quad (37a)$$

whereas if (H, π) is contained in one of the sets $\mathcal{C}^{i,j,k}$, then by the definition in line 23 we have

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \geq d_\sigma^i. \quad (37b)$$

Applying Lemma 21 using (36) shows that

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \leq d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) . \quad (38)$$

We will distinguish the cases where the inequality (38) is strict,

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) < d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) , \quad (39a)$$

and where it is tight,

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) . \quad (39b)$$

Altogether we distinguish four cases: whether (H, π) is contained in one of the sets $\mathcal{C}^{i,j}$ or $\mathcal{C}^{i,j,k}$, and whether the inequality (38) is strict or tight.

- $(H, \pi) \in \mathcal{C}^{i,j}$ and inequality (38) is strict. Combining (37a) and (39a) yields

$$d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) > d_\sigma^i . \quad (40)$$

By the definition of d_σ^i in line 8, it follows from (40) that if $(J, \pi|_J)$ was not already contained in \mathcal{H}_σ at the beginning of the repeat-loop (**) (in the i -th iteration of the repeat-loop (*)), then at this point $(J \setminus v_1, \pi|_{J \setminus v_1})$ was not contained in \mathcal{H}_σ either. By (35) there must then be some $j' < j$ such that $(J \setminus v_1, \pi|_{J \setminus v_1})$ was added to \mathcal{H}_σ in the j' -th iteration of the repeat-loop (**). Combining the first part of Lemma 17 and (40) shows that also $(J, \pi|_J)$ was added to \mathcal{H}_σ in the j' -th iteration of the repeat-loop (**). In any case $(J, \pi|_J)$ is already contained in \mathcal{H}_σ when (H, π) is added to this set. Applying the first part of Lemma 20 using (39a) yields that $w_{(H, \pi, \sigma)}(v_1) \leq w_{(J, \pi|_J, \sigma)}(v_1)$, completing the inductive step in this case.

- $(H, \pi) \in \mathcal{C}^{i,j,k}$ and inequality (38) is strict. Combining (37b) and (39a) yields $d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) > d_\sigma^i$. Therefore, if $(J, \pi|_J)$ is not already contained in \mathcal{H}_σ when (H, π) is added to this set via the set $\mathcal{C}^{i,j,k}$, then $(J, \pi|_J)$ is contained in $\mathcal{C}^{i,j,k}$ as well (recall (35) and the definition in line 23 and note that $(J, \pi|_J)$ satisfies the first condition in line 26), and added to \mathcal{H}_σ together with (H, π) . To complete the inductive step apply again the first part of Lemma 20 using (39a).
- $(H, \pi) \in \mathcal{C}^{i,j}$ and inequality (38) is tight. Combining (37a) and (39b) yields $d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) = d_\sigma^i$. Therefore, if $(J, \pi|_J)$ is not already contained in \mathcal{H}_σ when (H, π) is added to this set via the set $\mathcal{C}^{i,j}$, then $(J, \pi|_J)$ is contained in $\mathcal{C}^{i,j}$ as well (recall (35) and the definition in line 14), and added to \mathcal{H}_σ together with (H, π) . The definitions in line 18, 28, 30 and 31 show that in any case

$$w_{(H, \pi, \sigma)}(v_1) \stackrel{(23)}{=} w_\sigma(H, \pi) = w^i \quad \text{and} \quad w_{(J, \pi|_J, \sigma)}(v_1) \stackrel{(23)}{=} w_\sigma(J, \pi|_J) = w^{\bar{i}} \quad (41)$$

for some $\bar{i} \leq i$ with $\alpha_{\bar{i}} = \sigma$. Applying the second part of Lemma 16 using (41) yields that $w_{(H, \pi, \sigma)}(v_1) \leq w_{(J, \pi|_J, \sigma)}(v_1)$, completing the inductive step in this case.

- $(H, \pi) \in \mathcal{C}^{i,j,k}$ and inequality (38) is tight. Combining (37b) and (39b) yields

$$d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) \geq d_\sigma^i . \quad (42)$$

Therefore, if $(J, \pi|_J)$ is not already contained in \mathcal{H}_σ when (H, π) is added to this set via the set $\mathcal{C}^{i,j,k}$, then $(J, \pi|_J)$ is contained in $\mathcal{T}^{i,j,k}$ as well (recall (35) and the definition in line 23). Suppose for the sake of contradiction that $(J, \pi|_J)$ was not transferred from $\mathcal{T}^{i,j,k}$ to $\mathcal{C}^{i,j,k}$, i.e., that it violated the condition in line 26. By (42) and the first condition in line 26 we have

$$d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) = d_\sigma^i , \quad (43)$$

and by the second condition in line 26 there is a subgraph $\bar{J} \subseteq J$ with $v_1 \in \bar{J}$ and

$$(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\sigma^i) . \quad (44)$$

Combining (37b), (38) and (43) shows that

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i . \quad (45)$$

Clearly \bar{J} is a subgraph of H that contains the youngest vertex v_1 , which combined with (44) and (45) contradicts the fact that (H, π) satisfies the condition in line 26 (only graphs satisfying this condition are transferred from $\mathcal{T}^{i,j,k}$ to $\mathcal{C}^{i,j,k}$). Hence the graph $(J, \pi|_J)$ is either already contained in \mathcal{H}_σ or added to this set together with (H, π) via the set $\mathcal{C}^{i,j,k}$ at the latest.

It remains to show that $w_{(H, \pi, \sigma)}(v_1) \leq w_{(J, \pi|_J, \sigma)}(v_1)$. Suppose for the sake of contradiction that this inequality is violated. Using (39b) and the second part of Lemma 20 this implies that $w_\sigma(H, \pi)$ is defined in line 31 with \hat{i} defined in line 28, and that $w_\sigma(J, \pi|_J)$ is defined either in line 18 or in line 31 with \hat{i} defined in line 30. In the following we show that none of those cases can occur, as in each case, similarly to above, the existence of a graph $\bar{J} \subseteq H$ with $v_1 \in \bar{J}$ and $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$ causes (H, π) to violate the condition in line 27 (thus causing

a contradiction). First consider the case that $w_\sigma(J, \pi|_J)$ is defined in line 18, i.e., $(J, \pi|_J)$ is an element of one of the sets $\mathcal{C}^{\bar{i}, \bar{j}}$ defined in some iteration $\bar{i} \leq i$. From the definition in line 14 we know that

$$d_\sigma^\bar{i} = d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}) \stackrel{(39b)}{=} d_\theta(H, v_1, w_{(H, \pi, \sigma)}) . \quad (46)$$

We distinguish the subcases $\bar{j} = 1$ and $\bar{j} \geq 2$. If $\bar{j} = 1$ then $(J, \pi|_J)$ was added to the set $\mathcal{C}_\sigma(d_\sigma^\bar{i})$ in line 16. Using (46) it follows that $(J, \pi|_J) \in \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$, showing that the graph $\bar{J} = J$ itself causes (H, π) to violate the condition in line 27. Similarly, if $\bar{j} \geq 2$, then by the second part of Lemma 17 there is a subgraph $\bar{J} \subseteq J$ with $v_1 \in \bar{J}$ and $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\sigma^\bar{i})$. Using again (46) it follows that $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$, causing a contradiction also in this case. Now consider the case that $w_\sigma(J, \pi|_J)$ is defined in line 31 with \hat{i} defined in line 30. Then by the condition in line 27 there is a subgraph $\bar{J} \subseteq J$ with $v_1 \in \bar{J}$ and $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\theta(J, v_1, w_{(J, \pi|_J, \sigma)}))$. Using (39b) it follows that $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_\sigma(d_\theta(H, v_1, w_{(H, \pi, \sigma)}))$, causing (H, π) to violate the condition in line 27. This completes the proof. \square

3.6. Graphs in $\mathcal{C}^{i,j,k}$ irrelevant for Builder. When proving Proposition 6 in Section 4 we develop a Builder strategy along the lines of the algorithm COMPUTEWEIGHTS(). The following lemma will be crucial in this: It shows that Builder does not need to enforce any ordered graph $(H, \pi) \in \mathcal{S}(F)$ that is added to one of the families \mathcal{H}_s via one of the sets $\mathcal{C}^{i,j,k}$, as for each such graph there is an alternative ordering $\pi' \in \Pi(V(H))$ of its vertices such that (H, π') has already been added to \mathcal{H}_s via the set $\mathcal{C}^{i,j}$ with weights that are at least as good for Builder.

Lemma 23 (Partners between $\mathcal{C}^{i,j}$ and $\mathcal{C}^{i,j,k}$). *Let $\sigma \in [r]$ and consider some iteration i of the repeat-loop (*) with $\alpha_i = \sigma$ of the algorithm COMPUTEWEIGHTS(). In every iteration $j \geq 1$ of the repeat-loop (**), the set $\mathcal{C}^{i,j}$ defined in line 14 and the sets $\mathcal{C}^{i,j,k}$, $k \geq 1$, defined in line 24 and 32 during each iteration of the repeat-loop (***), satisfy the following: For any graph $(H, \pi) \in \mathcal{C}^{i,j,k}$, $\pi = (v_1, \dots, v_h)$, the graph (H, π') , defined by $\pi' := (v_{k+1}, v_1, v_2, \dots, v_k, v_{k+2}, \dots, v_h)$, is contained in $\mathcal{C}^{i,j}$ and satisfies $w_{(H, \pi, \sigma)}(u) \leq w_{(H, \pi', \sigma)}(u)$ for all $u \in H$.*

Proof. For the reader's convenience, Figure 3 illustrates the notations used throughout the proof.

We shall prove the following more technical claim: Let $\sigma \in [r]$ and consider some iteration i of the repeat-loop (*) with $\alpha_i = \sigma$ and some iteration j of the repeat-loop (**). For $k = 0$ and any graph $(H, \pi) \in \mathcal{C}^{i,j}$, $\pi = (v_1, \dots, v_h)$, and for $k \geq 1$ and any graph $(H, \pi) \in \mathcal{C}^{i,j,k}$, $\pi = (v_1, \dots, v_h)$, the graph (H, π') , defined by $\pi' := (v_{k+1}, v_1, v_2, \dots, v_k, v_{k+2}, \dots, v_h)$, is contained in $\mathcal{C}^{i,j}$ and satisfies $w_{(H, \pi, \sigma)}(u) \leq w_{(H, \pi', \sigma)}(u)$ for all $u \in H$.

We will argue at the end of the proof that any graph (H, π) contained in one of the sets $\mathcal{C}^{i,j,k}$ has at least three vertices, ensuring that all subgraphs used in the following arguments have at least one vertex.

To prove the auxiliary claim we consider a fixed iteration $j \geq 1$ of the repeat-loop (**) and argue by induction over k , the number of iterations of the repeat-loop (***). We choose the state before the beginning of the first iteration ($k = 0$) as our induction base. In this case $\pi' = \pi$ and the claim is trivially true.

For the induction step consider a graph $(H, \pi) \in \mathcal{C}^{i,j,k}$, $\pi = (v_1, \dots, v_h)$, that is added to \mathcal{H}_σ in the k -th iteration of the repeat-loop (***). By the definition of $\mathcal{T}^{i,j,k}$ in line 23 (recall that $\mathcal{C}^{i,j,k} \subseteq \mathcal{T}^{i,j,k}$) we clearly have

$$d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \geq d_\sigma^i . \quad (47)$$

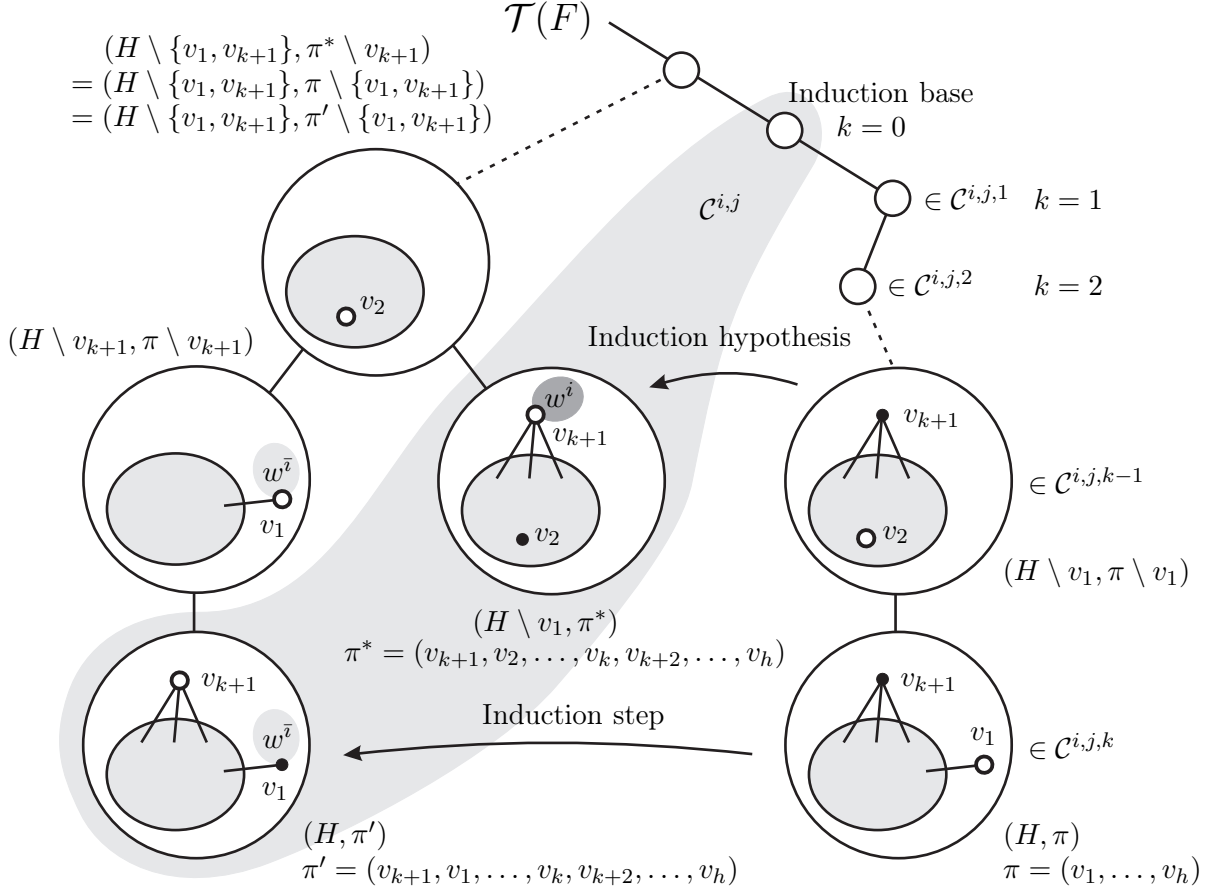


FIGURE 3. Notations used in the proof of Lemma 23. For each ordered graph, the respective youngest vertex is emphasized.

As (H, π) is obtained from $(H \setminus v_1, \pi \setminus v_1)$ by adding v_1 as the youngest vertex and all edges incident to it, we have

$$w_{(H, \pi, \sigma)}(u) = w_{(H \setminus v_1, \pi \setminus v_1, \sigma)}(u) \quad \text{for all } u \in H \setminus v_1. \quad (48)$$

If $j = 1$, the definition of d_σ^i in line 8 ensures that at the beginning of the j -th iteration of the repeat-loop (***) we have $d_\theta(J, u_1, w_{(J, \tau, \sigma)}) \leq d_\sigma^i$ for all $(J, \tau) \in \mathcal{C}(\mathcal{H}_\sigma, F)$, $\tau = (u_1, \dots, u_c)$. If $j > 1$, the same statement is true by the first part of Lemma 17. Thus if the inequality (47) is strict, then (H, π) was not in $\mathcal{C}(\mathcal{H}_\sigma, F)$ at the beginning of the j -th iteration of the repeat-loop (***), and thus $(H \setminus v_1, \pi \setminus v_1)$ was not in \mathcal{H}_σ at this point. If $k = 1$ this means that $(H \setminus v_1, \pi \setminus v_1)$ must have been added to \mathcal{H}_σ via the set $\mathcal{C}^{i,j}$. The same conclusion holds if $k = 1$ and the inequality (47) is tight, as otherwise (H, π) would have qualified for inclusion in $\mathcal{C}^{i,j}$.

Note that the conditions for inclusion into $\mathcal{T}^{i,j,k}$ and $\mathcal{C}^{i,j,k}$ in lines 23 and 26 do not change during the entire repeat-loop (***) except for the requirement that (H, π) is in the current set $\mathcal{C}(\mathcal{H}_\sigma, F)$. Thus if $k \geq 2$ then $(H \setminus v_1, \pi \setminus v_1)$ must have been added to \mathcal{H}_σ via $\mathcal{C}^{i,j,k-1}$ in the $(k-1)$ -th iteration of the repeat-loop (***)

In all cases we can apply the induction hypothesis and conclude that the graph $(H \setminus v_1, \pi^*)$, defined by $\pi^* := (v_{k+1}, v_2, v_3, \dots, v_k, v_{k+2}, \dots, v_h)$ (to be understood as $\pi \setminus v_1$ if $k = 1$) is contained in $\mathcal{C}^{i,j}$ and satisfies

$$w_{(H \setminus v_1, \pi \setminus v_1, \sigma)}(u) \leq w_{(H \setminus v_1, \pi^*, \sigma)}(u) \quad \text{for all } u \in H \setminus v_1. \quad (49)$$

By the definition of $\mathcal{C}^{i,j}$ in line 14, $(H \setminus v_1, \pi^*)$ satisfies

$$d_\theta(H \setminus v_1, v_{k+1}, w_{(H \setminus v_1, \pi^*, \sigma)}) = d_\sigma^i, \quad (50)$$

and the weight of its youngest vertex v_{k+1} is set to

$$w_{(H \setminus v_1, \pi^*, \sigma)}(v_{k+1}) \stackrel{(23)}{=} w_\sigma(H \setminus v_1, \pi^*) = w^i \quad (51)$$

in line 18.

Consider now the graph $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ and observe that it is a subgraph of (H, π) . Applying Lemma 21 and Lemma 22 hence yields

$$d_\theta(H \setminus v_{k+1}, v_1, w_{(H \setminus v_{k+1}, \pi \setminus v_{k+1}, \sigma)}) \geq d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \stackrel{(47)}{\geq} d_\sigma^i. \quad (52)$$

As $(H \setminus v_1, \pi^*) \in \mathcal{C}^{i,j}$ we must have had

$$(H \setminus \{v_1, v_{k+1}\}, \pi^* \setminus v_{k+1}) = (H \setminus \{v_1, v_{k+1}\}, \pi \setminus \{v_1, v_{k+1}\}) \in \mathcal{H}_\sigma \quad (53)$$

at the beginning of the j -th iteration of the repeat-loop (**).

As our next intermediate step we will show that the graph $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ (which is obtained from the graph in (53) by adding v_1 as the youngest vertex and all edges incident to it) is contained in \mathcal{H}_σ at the beginning of the j -th iteration of the repeat-loop (**) as well. We first consider the case that one of the inequalities in (52) is strict, i.e., we have

$$d_\theta(H \setminus v_{k+1}, v_1, w_{(H \setminus v_{k+1}, \pi \setminus v_{k+1}, \sigma)}) > d_\sigma^i. \quad (54)$$

If $j = 1$, it follows from (53) and (54) that $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ is indeed contained in \mathcal{H}_σ at the beginning of the j -th iteration of the repeat-loop (**), as otherwise we would obtain a contradiction to the definition of d_σ^i in line 8. The same conclusion holds for $j > 1$ by using (53), (54) and the first part of Lemma 17. Now consider the case that all inequalities in (52) are tight, i.e., we have

$$d_\theta(H \setminus v_{k+1}, v_1, w_{(H \setminus v_{k+1}, \pi \setminus v_{k+1}, \sigma)}) = d_\theta(H, v_1, w_{(H, \pi, \sigma)}) = d_\sigma^i. \quad (55)$$

Suppose that $j = 1$ and that $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ was not already contained in \mathcal{H}_σ at the beginning of the first iteration of the repeat-loop (**). Then by (53) and (55), this graph would be added to $\mathcal{C}^{i,1} = \mathcal{C}_\sigma(d_\sigma^i)$ in line 14. As $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ is a subgraph of (H, π) that contains the vertex v_1 , this observation together with the second equality in (55) contradicts the fact (H, π) satisfies the condition in line 26. The remaining subcase $j > 1$ can be proven analogously by using the second part of Lemma 17.

Therefore, we indeed have

$$(H \setminus v_{k+1}, \pi \setminus v_{k+1}) \in \mathcal{H}_\sigma \quad (56)$$

at the beginning of the j -th iteration of the repeat-loop (**). The weight assigned to the youngest vertex v_1 of this graph is

$$w_{(H \setminus v_{k+1}, \pi \setminus v_{k+1}, \sigma)}(v_1) = w^{\bar{i}} \geq w^i \quad (57)$$

for some $\bar{i} \leq i$ with $\alpha_{\bar{i}} = \sigma$, where the last inequality follows from the second part of Lemma 16.

We will show that the graph (H, π') , defined by $\pi' := (v_{k+1}, v_1, \dots, v_k, v_{k+2}, \dots, v_h)$ (which is obtained from $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ by adding v_{k+1} as the youngest vertex and all edges incident to it), satisfies the inductive claim: We first demonstrate that this graph is contained in $\mathcal{C}^{i,j}$ and then prove the claimed inequality between the vertex weights of (H, π) and (H, π') .

As a first step towards this goal we show that $d_\theta(H, v_{k+1}, w_{(H, \pi', \sigma)}) = d_\sigma^i$.

Clearly we have $(H \setminus \{v_1, v_{k+1}\}, \pi^* \setminus v_{k+1}) = (H \setminus \{v_1, v_{k+1}\}, \pi' \setminus \{v_1, v_{k+1}\})$, implying that

$$w_{(H \setminus v_1, \pi^*, \sigma)}(u) = w_{(H, \pi', \sigma)}(u) \quad \text{for all } u \in H \setminus \{v_1, v_{k+1}\}. \quad (58)$$

As (H, π') is obtained from $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ by adding v_{k+1} as the youngest vertex and all edges incident to it, we have

$$w_{(H, \pi', \sigma)}(v_1) = w_{(H \setminus v_{k+1}, \pi \setminus v_{k+1}, \sigma)}(v_1) \stackrel{(57)}{\geq} w^i. \quad (59)$$

Observe that (H, π') is a supergraph of $(H \setminus v_1, \pi^*)$, so applying Lemma 21 and Lemma 22 yields

$$d_\theta(H, v_{k+1}, w_{(H, \pi', \sigma)}) \leq d_\theta(H \setminus v_1, v_{k+1}, w_{(H \setminus v_1, \pi^*, \sigma)}) \stackrel{(50)}{=} d_\sigma^i. \quad (60)$$

Now suppose for the sake of contradiction that the first inequality in (60) is strict, i.e., we have

$$d_\theta(H, v_{k+1}, w_{(H, \pi', \sigma)}) < d_\sigma^i. \quad (61)$$

By (18) and (58) this implies that some

$$J' \in \arg \min_{J \subseteq H: v_{k+1} \in J} \left(\sum_{u \in J \setminus v_{k+1}} (1 + w_{(H, \pi', \sigma)}(u)) - e(J) \cdot \theta \right) \quad (62)$$

includes the vertex v_1 . But then we would have

$$\begin{aligned} d_\theta(H, v_1, w_{(H, \pi, \sigma)}) &\stackrel{(18)}{\leq} \sum_{u \in J' \setminus v_1} (1 + w_{(H, \pi, \sigma)}(u)) - e(J') \cdot \theta \\ &\stackrel{(48), (49)}{\leq} \sum_{u \in J' \setminus v_1} (1 + w_{(H \setminus v_1, \pi^*, \sigma)}(u)) - e(J') \cdot \theta \\ &= \sum_{u \in J' \setminus \{v_1, v_{k+1}\}} (1 + w_{(H \setminus v_1, \pi^*, \sigma)}(u)) + \underbrace{(1 + w_{(H \setminus v_1, \pi^*, \sigma)}(v_{k+1}))}_{\stackrel{(51)}{=} w^i} - e(J') \cdot \theta \\ &\stackrel{(58), (59)}{\leq} \sum_{u \in J' \setminus v_{k+1}} (1 + w_{(H, \pi', \sigma)}(u)) - e(J') \cdot \theta \\ &\stackrel{(18), (62)}{=} d_\theta(H, v_{k+1}, w_{(H, \pi', \sigma)}) \stackrel{(61)}{<} d_\sigma^i, \end{aligned}$$

contradicting (47). Hence (60) holds with equality and we have indeed

$$d_\theta(H, v_{k+1}, w_{(H, \pi', \sigma)}) = d_\sigma^i. \quad (63)$$

Clearly, as $(H \setminus v_1, \pi^*)$ is contained in $\mathcal{C}^{i,j}$, it follows that this graph is not contained in \mathcal{H}_σ at the beginning of the j -th iteration of the repeat-loop (*). As (H, π') is a supergraph of $(H \setminus v_1, \pi^*)$, it follows from Lemma 22 that (H, π') is not contained in \mathcal{H}_σ at this point either. Hence, by (56), (63) and the definition of $\mathcal{C}^{i,j}$ in line 14 we have $(H, \pi') \in \mathcal{C}^{i,j}$.

It remains to check the claimed inequality between the vertex weights of (H, π) and (H, π') . Note that $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ is a subgraph of (H, π) . We can hence apply Lemma 22 and, observing that (H, π') is obtained from $(H \setminus v_{k+1}, \pi \setminus v_{k+1})$ by adding v_{k+1} as the youngest vertex and all edges incident to it, obtain the desired inequality for all vertices in the set $\{v_1, \dots, v_h\} \setminus \{v_{k+1}\}$. For the vertex v_{k+1} , note that (48), (49) and (51) together yield $w_{(H, \pi, \sigma)}(v_{k+1}) \leq w^i$ and that $w_{(H, \pi', \sigma)}(v_{k+1}) \stackrel{(23)}{=} w_\sigma(H, \pi') = w^i$ by the definition in line 18.

This completes the inductive proof of Lemma 23.

It remains to show that every graph $(H, \pi) \in \mathcal{C}^{i,j,k}$ has at least three vertices, and that therefore all graphs used in the above arguments are well-defined and have at least one vertex. We show that all graphs $(H, \pi) \in \mathcal{S}(F)$ on at most two vertices that are ever added to the family \mathcal{H}_σ in the course of the algorithm are added to it via one of the sets $\mathcal{C}^{i,j}$ defined in line 14: As argued in the proof of Lemma 12 on page 24, this is true for the graph $(K_1, (v_1))$ (an isolated vertex), which is added

via the set $\mathcal{C}^{i,1}$ in the first iteration i for which $\alpha_i = \sigma$. Once we have $\mathcal{H}_\sigma = \{(K_1, (v_1))\}$, the only two ordered subgraphs of F on two vertices, a single edge and two isolated vertices, are contained in $\mathcal{C}(\mathcal{H}_\sigma, F)$. Using this fact together with the observation that $(K_1, (v_1))$ is contained in $\mathcal{C}_\sigma(0)$ and is a subgraph of both of them, it is easy to check that each of those two graphs can only be added to \mathcal{H}_σ via one of the sets $\mathcal{C}^{i,j}$: If the $d_\theta()$ -value of one of these graphs is equal to 0, then it is added via $\mathcal{C}^{i,2}$. Otherwise its $d_\theta()$ -value is strictly smaller than 0 and it is added via $\mathcal{C}^{i,1}$ for some $i > i$. \square

3.7. Further properties of the algorithm. The next lemma implies in particular that for a graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$ and a subgraph $J \subseteq H$, $v_1 \in J$, minimizing the right hand side of the definition of $d_\theta(H, v_1, w_{(H, \pi, s)})$ in (18), the inequality stated in Lemma 22 is in fact an equality. As a consequence, in all situations where the vertex weights of a subgraph $J \subseteq H$ are relevant, these weights only depend on $(J, \pi|_J)$ and not on the ‘context’ H . This is far from clear a priori, and in fact not true for arbitrary subgraphs $J \subseteq H$.

Lemma 24 (Irrelevant context of $d_\theta()$ -minimizing subgraphs). *Throughout the algorithm COMPUTE-WEIGHTS(), for every $s \in [r]$, any graph $(H, \pi) \in \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$, $\pi = (v_1, \dots, v_h)$, and any graph \hat{J} from the family*

$$\arg \min_{J \subseteq H: v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(H, \pi, s)}(u)) - e(J) \cdot \theta \right) \quad (64)$$

we have

$$w_{(H, \pi, s)}(u) = w_{(\hat{J}, \pi|_{\hat{J}}, s)}(u) \quad \text{for all } u \in \hat{J} \setminus v_1. \quad (65a)$$

Moreover, if $(H, \pi) \in \mathcal{H}_s$, then we have

$$w_{(H, \pi, s)}(v_1) = w_{(\hat{J}, \pi|_{\hat{J}}, s)}(v_1). \quad (65b)$$

Note that by Lemma 15 and Lemma 22, all vertex weights referred to in the formulation of Lemma 24 are finite values. We will not mention this explicitly again in the following.

The following two auxiliary statements are only used for proving Lemma 24.

Lemma 25. *Let H be a graph with $V(H) = \{v_1, \dots, v_h\}$, $w : V(H) \setminus \{v_1\} \rightarrow \mathbb{R}$ an arbitrary weight function, $\theta > 0$ a real number, and let \hat{J} be a graph from the family*

$$\arg \min_{J \subseteq H: v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w(u)) - e(J) \cdot \theta \right). \quad (66)$$

Moreover, let v_k be a vertex contained in \hat{J} and \tilde{J} a graph from the family

$$\arg \min_{J \subseteq H \setminus \{v_1, \dots, v_{k-1}\}: v_k \in J} \left(\sum_{u \in J \setminus v_k} (1 + w(u)) - e(J) \cdot \theta \right). \quad (67)$$

Then the graph $\tilde{J} \cap \hat{J}$ is also contained in the family (67). In particular, for two graphs J', J'' from the family (66) the graph $J' \cap J''$ is also contained in (66).

Proof. In order to simplify notation, we introduce for a real number $\theta > 0$, any graph H , any vertex $v \in H$ and any weight function $w : V(H) \setminus \{v\} \rightarrow \mathbb{R}$, the function

$$\lambda_\theta^-(H, v, w) := \sum_{u \in H \setminus v} (1 + w(u)) - e(H) \cdot \theta. \quad (68)$$

As for the definition of $d_\theta()$ in (18), it is also convenient here to allow functions w in the third argument of $\lambda_\theta^-()$ whose domain is strictly larger than the set $V(H) \setminus \{v\}$. Of course, for the value of $\lambda_\theta^-(H, v, w)$ only the values $w(u)$ for all $u \in H \setminus v$ are relevant.

By the choice of \tilde{J} in (67) and by (68), we have

$$\lambda_{\theta}^{-}(\tilde{J}, v_k, w) \leq \lambda_{\theta}^{-}(\tilde{J} \cap \hat{J}, v_k, w) . \quad (69)$$

This inequality, however, must be tight, as otherwise the second inequality in

$$\lambda_{\theta}^{-}(\hat{J} \cup \tilde{J}, v_1, w) \stackrel{(68)}{=} \lambda_{\theta}^{-}(\hat{J}, v_1, w) + \lambda_{\theta}^{-}(\tilde{J}, v_k, w) - \lambda_{\theta}^{-}(\tilde{J} \cap \hat{J}, v_k, w) \stackrel{(69)}{\leq} \lambda_{\theta}^{-}(\hat{J}, v_1, w)$$

would be strict, contradicting the choice of \hat{J} in (66). This proves the lemma. \square

The next auxiliary statement will be used to prove Lemma 24 by induction.

Lemma 26. *The following invariant holds throughout the algorithm COMPUTEWEIGHTS(). Let $s \in [r]$ and let (H, π) , $\pi = (v_1, \dots, v_h)$, be a graph in \mathcal{H}_s , and suppose that every graph J' from the family*

$$\arg \min_{J \subseteq H: v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(H, \pi, s)}(u)) - e(J) \cdot \theta \right) \quad (70)$$

satisfies

$$w_{(H, \pi, s)}(u) = w_{(J', \pi|_{J'}, s)}(u) \quad \text{for all } u \in J' \setminus v_1 . \quad (71)$$

Then every such graph J' satisfies

$$w_{(H, \pi, s)}(v_1) = w_{(J', \pi|_{J'}, s)}(v_1) . \quad (72)$$

Proof. Let J' be a graph from the family (70) and note that

$$d_{\theta}(H, v_1, w_{(H, \pi, s)}) \stackrel{(18), (70)}{=} d_{\theta}(J', v_1, w_{(H, \pi, s)}) \stackrel{(71)}{=} d_{\theta}(J', v_1, w_{(J', \pi|_{J'}, s)}) . \quad (73)$$

By Lemma 22 we clearly have $w_{(H, \pi, s)}(v_1) \leq w_{(J', \pi|_{J'}, s)}(v_1)$. Consequently, using (73) and applying the second part of Lemma 20, the only way that $w_{(H, \pi, s)}(v_1)$ can be different from $w_{(J', \pi|_{J'}, s)}(v_1)$ is if $w_s(H, \pi)$ is defined either in line 18 or in line 31 with \hat{i} defined in line 30, and $w_s(J', \pi|_{J'})$ is defined in line 31 with \hat{i} defined in line 28. We will show that none of those cases can occur.

- First consider the case that $w_s(H, \pi)$ is defined in line 18 in some iteration i of the repeat-loop (*) (for which $\alpha_i = s$) and the first iteration $j = 1$ of the repeat-loop (**), i.e., (H, π) is contained in $\mathcal{C}^{i,1}$ and satisfies $d_{\theta}(H, v_1, w_{(H, \pi, s)}) = d_s^i$ and $w_s(H, \pi) = w^i$. Then by (73) and by Lemma 18 and Lemma 22 the graph $(J', \pi|_{J'})$ must be contained in $\mathcal{C}^{i,1}$ as well and is added to \mathcal{H}_{σ} together with (H, π) . Hence we have $w_s(J', \pi|_{J'}) = w^i$ by the definition in line 18, proving (72) in this case.
- Now consider the case that $w_s(H, \pi)$ is defined in line 18 in some iteration i of the repeat-loop (*) (for which $\alpha_i = s$) and some iteration $j > 1$ of the repeat-loop (**), i.e., (H, π) is contained in $\mathcal{C}^{i,j}$ and satisfies

$$d_{\theta}(H, v_1, w_{(H, \pi, s)}) = d_s^i . \quad (74)$$

Then (H, π) must have been in $\mathcal{C}(\mathcal{H}_s, F)$ at the end of the previous iteration of the repeat-loop (**), and by the second part of Lemma 17 there is a subgraph $\bar{J} \subseteq H$ with $v_1 \in \bar{J}$ and $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_s(d_s^i)$. From the definitions in line 14 and line 16 it follows that

$$d_{\theta}(\bar{J}, v_1, w_{(\bar{J}, \pi|_{\bar{J}}, s)}) = d_s^i . \quad (75)$$

Fix some graph J'' from the family

$$\arg \min_{J \subseteq \bar{J}: v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(\bar{J}, \pi|_{\bar{J}}, s)}(u)) - e(J) \cdot \theta \right) \quad (76)$$

and note that

$$d_\theta(\bar{J}, v_1, w_{(\bar{J}, \pi|_{\bar{J}, s})}) \stackrel{(18), (76)}{=} \sum_{u \in J'' \setminus v_1} (1 + w_{(\bar{J}, \pi|_{\bar{J}, s})}(u)) - e(J'') \cdot \theta . \quad (77)$$

By Lemma 22 we have

$$w_{(H, \pi, s)}(u) \leq w_{(\bar{J}, \pi|_{\bar{J}, s})}(u) \quad \text{for all } u \in \bar{J} . \quad (78)$$

We hence have

$$d_\theta(H, v_1, w_{(H, \pi, s)}) \stackrel{(18)}{\leq} \sum_{u \in J'' \setminus v_1} (1 + w_{(H, \pi, s)}(u)) - e(J'') \cdot \theta \stackrel{(77), (78)}{\leq} d_\theta(\bar{J}, v_1, w_{(\bar{J}, \pi|_{\bar{J}, s})}) ,$$

which combined with (74) and (75) shows that the graph $(J'', \pi|_{J''})$ is contained in the family (70). Applying Lemma 25 yields that the graph $(J' \cap J'', \pi|_{J' \cap J''})$ is also contained in the family (70). Analogously to (73) we have

$$d_\theta(H, v_1, w_{(H, \pi, s)}) = d_\theta(J' \cap J'', v_1, w_{(J' \cap J'', \pi|_{J' \cap J'', s})}) ,$$

which combined with (74) shows that

$$d_\theta(J' \cap J'', v_1, w_{(J' \cap J'', \pi|_{J' \cap J'', s})}) = d_s^i . \quad (79)$$

Clearly $J' \cap J''$ is a subgraph of \bar{J} that contains the youngest vertex v_1 . Using this observation and (79) and applying Lemma 18 and Lemma 22, the fact that $(\bar{J}, \pi|_{\bar{J}})$ is contained in $\mathcal{C}_s(d_s^i)$ implies that $(J' \cap J'', \pi|_{J' \cap J''})$ is contained in $\mathcal{C}_s(d_s^i)$ as well. But as $J' \cap J''$ is also a subgraph of J' (that contains the youngest vertex v_1), this implies with (73) and (74) that the graph $(J', \pi|_{J'})$ violates the condition in line 27, which yields the desired contradiction.

- Finally consider the case that $w_s(H, \pi)$ is defined in line 31 (in some iteration i of the repeat-loop (*)) with \hat{i} defined in line 30. By the conditions in line 26 and line 27 we have $d_\theta(H, v_1, w_{(H, \pi, s)}) > d_s^i$ and there is a graph $\bar{J} \subseteq H$ with $v_1 \in \bar{J}$ and $(\bar{J}, \pi|_{\bar{J}}) \in \mathcal{C}_s(d_\theta(H, v_1, w_{(H, \pi, s)}))$. Using the definitions in line 14 and line 16, as well as the first part of Lemma 16, it follows that

$$d_\theta(H, v_1, w_{(H, \pi, s)}) = d_s^{\bar{i}}$$

and

$$d_\theta(\bar{J}, v_1, w_{(\bar{J}, \pi|_{\bar{J}, s})}) = d_s^{\bar{i}}$$

for some $\bar{i} < i$. From here the proof continues analogously to the second case (where d_s^i needs to be replaced by $d_s^{\bar{i}}$), concluding that $(J', \pi|_{J'})$ must violate the condition in line 27, which again yields the desired contradiction. \square

Proof of Lemma 24. We argue by induction over the number of vertices of H . The claim clearly holds if H consists only of a single vertex, as then $\hat{J} = H$ is the only graph contained in the family (64). This settles the base of the induction.

For the induction step suppose that H has at least two vertices, and let \hat{J} be a graph from the family (64). Clearly, $(H \setminus v_1, \pi \setminus v_1)$ is contained in \mathcal{H}_s (recall the definition of $\mathcal{C}()$ in (22)). Applying Lemma 22 we obtain that $(\hat{J} \setminus v_1, \pi|_{\hat{J} \setminus v_1})$ is contained in \mathcal{H}_s as well and that

$$w_{(H, \pi, s)}(v_i) \leq w_{(\hat{J}, \pi|_{\hat{J}, s})}(v_i) \quad \text{for all } v_i \in \hat{J} \setminus v_1 . \quad (80)$$

We will first show that this inequality is tight for all $v_i \in \hat{J} \setminus v_1$ (which is exactly the statement of (65a)). Suppose for the sake of contradiction that the inequality in (80) is strict for some $v_i \in \hat{J} \setminus v_1$,

and choose the largest index k for which this is the case, i.e.

$$w_{(H,\pi,s)}(v_k) < w_{(\hat{J},\pi|_{\hat{J}},s)}(v_k) \quad (81)$$

and

$$w_{(H,\pi,s)}(v_i) = w_{(\hat{J},\pi|_{\hat{J}},s)}(v_i) \quad \text{for all } v_i \in \hat{J} \setminus \{v_1, \dots, v_k\} \quad (82)$$

(we clearly have $k \geq 2$). Fix a graph \tilde{J} from the family

$$\arg \min_{J \subseteq H \setminus \{v_1, \dots, v_{k-1}\} : v_k \in J} \left(\sum_{u \in J \setminus v_k} (1 + w_{(H,\pi,s)}(u)) - e(J) \cdot \theta \right) \quad (83)$$

and observe that by Lemma 25, also the graph $\tilde{J} \cap \hat{J}$ is contained in the family (83). By (82) the same graph $\tilde{J} \cap \hat{J}$ is also contained in the family

$$\arg \min_{J \subseteq \hat{J} \setminus \{v_1, \dots, v_{k-1}\} : v_k \in J} \left(\sum_{u \in J \setminus v_k} (1 + w_{(\hat{J},\pi|_{\hat{J}},s)}(u)) - e(J) \cdot \theta \right) .$$

By induction, we therefore have $w_{(\tilde{J} \cap \hat{J}, \pi|_{\tilde{J} \cap \hat{J}}, s)}(v_k) = w_{(H,\pi,s)}(v_k)$ and $w_{(\tilde{J} \cap \hat{J}, \pi|_{\tilde{J} \cap \hat{J}}, s)}(v_k) = w_{(\hat{J}, \pi|_{\hat{J}}, s)}(v_k)$, which together contradicts (81) and shows that (80) holds with equality for all $v_i \in \hat{J} \setminus v_1$, thus proving (65a).

The relation (65b) follows from (65a) by applying Lemma 26. \square

Lemma 24 allows us to derive the next statement, which is similar in spirit but considers $\lambda_\theta()$ -values instead of $d_\theta()$ -values.

Lemma 27 (Irrelevant context of $\lambda_\theta()$ -minimizing subgraphs). *For every $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, we have*

$$\min_{J \subseteq H} \lambda_\theta(J, w_{(H,\pi,s)}) = \min_{J \subseteq H} \lambda_\theta(J, w_{(J,\pi|_J,s)}) .$$

Proof. Using the definition of $\lambda_\theta()$ in (19), we obtain from Lemma 22 that

$$\min_{J \subseteq H : v_1 \in J} \lambda_\theta(J, w_{(H,\pi,s)}) \leq \min_{J \subseteq H : v_1 \in J} \lambda_\theta(J, w_{(J,\pi|_J,s)}) . \quad (84)$$

From

$$\min_{J \subseteq H : v_1 \in J} \lambda_\theta(J, w_{(H,\pi,s)}) \stackrel{(19)}{=} \min_{J \subseteq H : v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(H,\pi,s)}(u)) - e(J) \cdot \theta \right) + 1 + w_{(H,\pi,s)}(v_1) \quad (85)$$

it follows that the minimum on the left hand side of (84) is attained for some graph

$$\hat{J} \in \arg \min_{J \subseteq H : v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(H,\pi,s)}(u)) - e(J) \cdot \theta \right) . \quad (86)$$

We can hence apply Lemma 24 and obtain

$$\min_{J \subseteq H : v_1 \in J} \lambda_\theta(J, w_{(H,\pi,s)}) \stackrel{(85),(86)}{=} \sum_{u \in \hat{J}} (1 + w_{(H,\pi,s)}(u)) - e(\hat{J}) \cdot \theta \stackrel{(19),(65)}{=} \lambda_\theta(\hat{J}, w_{(\hat{J},\pi|_{\hat{J}},s)}) ,$$

which shows that the inequality in (84) is tight. The lemma now follows by combining the resulting identity with the identity

$$\min_{J \subseteq H} \lambda_\theta(J, w_{(H,\pi,s)}) = \min_{\substack{1 \leq i \leq h \\ J \subseteq H \setminus \{v_1, \dots, v_{i-1}\} : v_i \in J}} \lambda_\theta(J, w_{(H,\pi,s)}) .$$

\square

We are now ready to state and prove the relation between $\Lambda_\theta(F, r)$ as defined in (24) and the parameter $\beta_i = 1 + \sum_{s \in [r]} d_s^i$ used in our informal explanation of the algorithm COMPUTEWEIGHTS() in Section 3.3.

Lemma 28 (Relation between $\Lambda_\theta(F, r)$ and d_s^i). *For any input sequence $\alpha \in [r]^{r \cdot |S(F)|}$ of the algorithm COMPUTEWEIGHTS() we have*

$$\max_{\substack{s \in [r] \\ \pi \in \Pi(V(F))}} \min_{H \subseteq F} \lambda_\theta(H, w_{(H, \pi|_H, s)}) = 1 + \sum_{s \in [r]} d_s^{\tilde{i}}, \quad (87)$$

where \tilde{i} is the smallest integer i for which $(F, \pi) \in \mathcal{C}^{i, j}$ for some $\pi \in \Pi(V(F))$ and some integer $j \geq 1$, and d_s^i and $\mathcal{C}^{i, j}$ are defined in line 8 and line 14.

Proof. Throughout the proof, we will repeatedly use that, as a consequence of the first part of Lemma 16, each of the values d_t^i , $t \in [r]$, is non-increasing with i , and that the sum

$$1 + \sum_{t \in [r]} d_t^i \quad (88)$$

is decreasing with i .

For every $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, we have

$$\min_{J \subseteq H: v_1 \in J} \lambda_\theta(J, w_{(H, \pi, s)}) \stackrel{(18), (19)}{=} d_\theta(H, v_1, w_{(H, \pi, s)}) + 1 + w_{(H, \pi, s)}(v_1). \quad (89)$$

If $(H, \pi) \in \mathcal{C}^{i, j}$ for some integers $i, j \geq 1$ with $\alpha_i = s$, then by using the definition in line 14 and by combining (23) with the definitions in line 10 and 18 we obtain from (89) that

$$\min_{J \subseteq H: v_1 \in J} \lambda_\theta(J, w_{(H, \pi, s)}) = 1 + \sum_{t \in [r]} d_t^i. \quad (90a)$$

Similarly, if $(H, \pi) \in \mathcal{C}^{i, j, k}$ for some integers $i, j, k \geq 1$ with $\alpha_i = s$, then by using the definition in line 23 (recall that $\mathcal{C}^{i, j, k} \subseteq \mathcal{T}^{i, j, k}$) and by combining (23) with the definitions in line 10 and 31 we obtain from (89), using the monotonicity of the values d_t^i in i , that

$$\min_{J \subseteq H: v_1 \in J} \lambda_\theta(J, w_{(H, \pi, s)}) \geq 1 + \sum_{t \in [r]} d_t^i. \quad (90b)$$

By Lemma 15, in the maximization in (87) it suffices to consider those $s \in [r]$ and vertex orderings $\pi \in \Pi(V(F))$, $\pi = (v_1, \dots, v_f)$, for which $(F, \pi) \in \mathcal{H}_s$. We clearly have

$$\min_{H \subseteq F} \lambda_\theta(H, w_{(F, \pi, s)}) = \min_{\substack{1 \leq c \leq f \\ H \subseteq F \setminus \{v_1, \dots, v_{c-1}\}: v_c \in H}} \lambda_\theta(H, w_{(F, \pi, s)}). \quad (91)$$

If $(F, \pi) \in \mathcal{C}^{i, j}$ for some integers $i, j \geq 1$ with $\alpha_i = s$, then by (90) and the monotonicity of the sum (88) in i , the minimum on the right hand side of (91) is attained for $c = 1$, yielding

$$\min_{H \subseteq F} \lambda_\theta(H, w_{(F, \pi, s)}) = 1 + \sum_{t \in [r]} d_t^i. \quad (92)$$

If $(F, \pi) \in \mathcal{C}^{i, j, k}$ for some integers $i, j, k \geq 1$ with $\alpha_i = s$, then by Lemma 23, the graph (F, π') , defined by $\pi' := (v_{k+1}, v_1, v_2, \dots, v_k, v_{k+2}, \dots, v_f)$, is contained in $\mathcal{C}^{i, j}$ and satisfies

$$w_{(F, \pi, s)}(u) \leq w_{(F, \pi', s)}(u) \quad \text{for all } u \in F, \quad (93)$$

implying that

$$\min_{H \subseteq F} \lambda_\theta(H, w_{(F, \pi, s)}) \stackrel{(93)}{\leq} \min_{H \subseteq F} \lambda_\theta(H, w_{(F, \pi', s)}) \stackrel{(92)}{=} 1 + \sum_{t \in [r]} d_t^i. \quad (94)$$

By Lemma 27 we can replace the weight function $w_{(F,\pi,s)}$ on the left hand side of (92) by $w_{(H,\pi|_H,s)}$ and the weight functions $w_{(F,\pi,s)}$ and $w_{(F,\pi',s)}$ in (94) by $w_{(H,\pi|_H,s)}$ and $w_{(H,\pi'|_H,s)}$, respectively. From the two modified equations the claim follows immediately, observing that as a consequence of the monotonicity of the sum (88) in i , their respective right hand sides are maximized for $i = \tilde{i}$ as defined in the lemma. \square

4. BUILDER IN THE DETERMINISTIC GAME

In this section we prove Proposition 6 by explicitly constructing, for F , r , θ and β as in the proposition, a Builder strategy that enforces a monochromatic copy of F in the deterministic game with r colors in at most a_{\max} steps (where $a_{\max} = a_{\max}(F, r)$ is defined in (97) below), and that respects the generalized density restriction (θ, β) .

4.1. The pigeonholing. We will derive Builder's strategy from the algorithm COMPUTEWEIGHTS() (Algorithm 1 on page 20) in two steps, using an abstract version of the game as an intermediate step. The reader should not be put off by our introducing yet another game — this abstract game is merely a convenience to separate a conceptually simple but important pigeonholing argument from the more interesting part of the proof. We give the pigeonholing argument in detail because there are some subtleties involved, and also because we want to derive an *explicit* upper bound $a_{\max} = a_{\max}(F, r)$ (that in particular does not depend on θ) on the number of steps that Builder needs to enforce a copy of F in the original deterministic game.

The abstract game is played by two players AbstractBuilder and AbstractPainter that correspond to the players of the original deterministic game. The state of the abstract game after t steps is described by a list (G^1, \dots, G^t) of r -colored graphs G^i , $1 \leq i \leq t$, where the same r -colored graph may appear several times in the list. (Intuitively, these entries represent r -colored graphs of which Builder can enforce isolated copies on the board of the actual deterministic game, where by an isolated copy of some (r -colored) graph G on the board we mean a copy of G that is the union of one or several components.) In each step $t + 1$ of the abstract game, AbstractBuilder constructs a new graph by choosing an arbitrary subset \mathcal{X} of the index set $\{1, \dots, t\}$, and connecting an additional vertex v in an arbitrary way to the disjoint union of the graphs G^i , $i \in \mathcal{X}$. AbstractPainter then chooses a color $s \in [r]$ for v , and the resulting r -colored graph G^{t+1} is added to the list. The game starts with the empty list (and thus the first graph G^1 constructed is simply an isolated vertex), and AbstractBuilder's goal is to create an r -colored graph G^t that contains a monochromatic copy of F . Similarly to before we say that an AbstractBuilder strategy satisfies the generalized density restriction (θ, β) for given values $\theta > 0$ and β if, at all times, all subgraphs H with $v(H) \geq 1$ of all graphs G^i in AbstractBuilder's list satisfy $\mu_\theta(H) \geq \beta$ (recall (7)).

The following lemma relates the abstract game to the original deterministic game.

Lemma 29 (Link between abstract and original deterministic game). *Let ABSTRACTSTRATEGY be an arbitrary AbstractBuilder strategy for the abstract game with r colors. If ABSTRACTSTRATEGY enforces a monochromatic copy of F in at most t_{\max} steps, then it gives rise to a Builder strategy STRATEGY for the original deterministic game with r colors that enforces a monochromatic copy of F in at most $(r+1)^{t_{\max}}$ steps. Furthermore, if ABSTRACTSTRATEGY satisfies the generalized density restriction (θ, β) for given values $\theta > 0$ and $\beta \geq 0$, then also STRATEGY satisfies the generalized density restriction (θ, β) .*

Proof. We simultaneously capture all possible ways the abstract game may evolve if AbstractBuilder plays according to ABSTRACTSTRATEGY by an r -ary rooted tree \mathcal{T} in which a node at depth t is a list $b = (G^1, \dots, G^t)$ of r -colored graphs G^i , $1 \leq i \leq t$, and has as its r children the nodes $b_s = (G^1, \dots, G^t, G_s^{t+1})$, $s \in [r]$, where G_s^{t+1} is obtained from G^1, \dots, G^t by applying the next

construction step of ABSTRACTSTRATEGY and coloring the new vertex with color s . Thus the graphs G_s^{t+1} differ only in the color assigned to the new vertex.

We assume w.l.o.g. that AbstractBuilder stops playing as soon as a monochromatic copy of F is created. Thus if $b = (G^1, \dots, G^t)$ is a leaf of \mathcal{T} , the graph G^t (the last graph constructed) contains a monochromatic copy of F . Furthermore, by the assumption of the lemma, the depth of the strategy tree \mathcal{T} is bounded by t_{\max} . In the following we assume w.l.o.g. that the depth of \mathcal{T} is exactly t_{\max} .

To derive STRATEGY from ABSTRACTSTRATEGY, we compute for each node $b = (G^1, \dots, G^t)$ of \mathcal{T} a function $f_b : \{G^1, \dots, G^t\} \rightarrow \mathbb{N}_0$ that specifies for each of the graphs G^i the number of isolated copies of G^i that are needed to implement the strategy ABSTRACTSTRATEGY in the original deterministic game. This can be done recursively as follows.

If $b = (G^1, \dots, G^t)$ is a leaf of \mathcal{T} , we set

$$f_b(G^t) := 1 \quad (95a)$$

and

$$f_b(G^i) := 0 \quad , \quad 1 \leq i \leq t-1 \quad . \quad (95b)$$

If $b = (G^1, \dots, G^t)$ is an internal node of \mathcal{T} , then letting $\mathcal{X}_b \subseteq \{1, \dots, t\}$ denote the index set of the graphs that are used in the construction step corresponding to b , and denoting the descendants of b by $b_s = (G^1, \dots, G^t, G_s^{t+1})$, $s \in [r]$, as before, we define for $1 \leq i \leq t$,

$$f_b(G^i) := \begin{cases} \max_{s \in [r]} f_{b_s}(G^i) \quad , & \text{if } i \notin \mathcal{X}_b \quad , \\ \max_{s \in [r]} f_{b_s}(G^i) + \sum_{s \in [r]} f_{b_s}(G_s^{t+1}) \quad , & \text{if } i \in \mathcal{X}_b \quad . \end{cases} \quad (95c)$$

With these definitions, STRATEGY is obtained from ABSTRACTSTRATEGY by proceeding as described by the strategy tree \mathcal{T} , and repeating every construction step corresponding to a given node b exactly $\sum_{s \in [r]} f_{b_s}(G_s^{t+1})$ times, each time connecting a new vertex to (previously unused) isolated copies of the graphs G^i , $i \in \mathcal{X}_b$, on the board as specified by the corresponding step of the abstract game. By the pigeonhole principle, this guarantees that regardless of how Painter plays there is a color σ such that at least $f_{b_\sigma}(G_\sigma^{t+1})$ isolated copies of G_σ^{t+1} are created, and by our recursive definition in (95c) it also follows that at least $f_{b_\sigma}(G_\sigma^i)$ isolated copies of each graph G_σ^i , $1 \leq i \leq t$, are left unused. Thus Builder may continue with the construction step corresponding to the node b_σ . This shows that at every node $b = (G^1, \dots, G^t)$ of \mathcal{T} , Builder has at least $f_b(G^i)$ isolated copies of every graph G^i available. In particular, when he reaches a leaf of \mathcal{T} , due to (95a) he will have created at least one copy of a graph G^t containing a monochromatic copy of F .

This shows that STRATEGY indeed creates a monochromatic copy of F in the original deterministic game, and it remains to bound the number of steps it needs to do so. For every $t = 1, \dots, t_{\max}$ we denote by c_t the maximum of $f_b(G^i)$ over all nodes $b = (G^1, \dots, G^t)$ at depth t in \mathcal{T} and all $1 \leq i \leq t$. It follows from (95c) that

$$c_t \leq (r+1)c_{t+1} \quad , \quad (96a)$$

and by (95a) and (95b) we have

$$c_{t_{\max}} = 1 \quad . \quad (96b)$$

By definition of the rule how often to repeat each step of ABSTRACTSTRATEGY in STRATEGY, the number of repetitions of a step that corresponds to a node b at depth t in \mathcal{T} is bounded by $r \cdot c_{t+1}$. It follows that the total number of Builder steps when executing STRATEGY is bounded by

$$\sum_{t=0}^{t_{\max}-1} r \cdot c_{t+1} \stackrel{(96)}{\leq} r \sum_{t=0}^{t_{\max}-1} (r+1)^t \leq (r+1)^{t_{\max}} \quad ,$$

as claimed.

Furthermore, as the strategy STRATEGY differs from ABSTRACTSTRATEGY merely in how often (Abstract)Builder's construction steps are repeated, and because for $\beta \geq 0$ it suffices to check the condition (7) for all *connected* subgraphs H of the board, it follows that with ABSTRACTSTRATEGY also STRATEGY satisfies the generalized density restriction (θ, β) . \square

4.2. Builder's strategy and proof of Proposition 6. We now present AbstractBuilder's strategy ABSTRACTBUILD(F, r, θ) that will yield our final Builder strategy BUILD(F, r, θ) via Lemma 29. Throughout this section, F , r , and θ are fixed, and we usually omit these arguments when we refer to ABSTRACTBUILD(F, r, θ) or COMPUTEWEIGHTS(F, r, θ, α).

The strategy ABSTRACTBUILD() proceeds in rounds along the lines of the algorithm COMPUTEWEIGHTS(). (As before, the term 'round' refers to one iteration of the repeat-loop (*) of COMPUTEWEIGHTS().) ABSTRACTBUILD() maintains, for each color $s \in [r]$, a family $\mathcal{G}_s \subseteq \mathcal{S}(F)$ and a mapping G_s from \mathcal{G}_s to the r -colored graphs in AbstractBuilder's list. For any $(H, \pi) \in \mathcal{G}_s$ the graph $G_s(H, \pi)$ will always contain a distinguished monochromatic copy of H in color s to which we will refer as the *central copy of H in $G_s(H, \pi)$* ; it is however possible that this copy was constructed in an order different from π (this is where we make crucial use of Lemma 23 proved in Section 3.6).

At the same time, ABSTRACTBUILD() extracts a sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ from AbstractPainter's coloring decisions such that the following holds: After each round, the families \mathcal{G}_s contain all graphs from the families \mathcal{H}_s occurring after the same number of rounds of COMPUTEWEIGHTS() with input sequence α . We will also see that for each graph $(H, \pi) \in \mathcal{H}_s$, the graph $G_s(H, \pi)$ on AbstractBuilder's list can indeed be used in further construction steps (without violating some given generalized density restriction) as indicated by the weight function $w_{(H, \pi, s)}$ computed by COMPUTEWEIGHTS() with input sequence α .

In order to construct a sequence α for which the above statements hold, ABSTRACTBUILD() uses variables defined by the algorithm COMPUTEWEIGHTS() for several different input sequences. We will use the following notations: For any sequence $\alpha \in [r]^{i-1}$ and any $s \in [r]$ we let $\alpha \circ s \in [r]^i$ denote the concatenation of α with s . When we refer to the algorithm COMPUTEWEIGHTS() with some input sequence $\alpha \in [r]^i$, we tacitly assume that α is extended arbitrarily to a sequence $\alpha' \in [r]^{r \cdot |\mathcal{S}(F)|}$ with prefix α . As we will only use this convention when we refer to variables defined in the first i iterations of the repeat-loop (*) of COMPUTEWEIGHTS(), the values of α' beyond the prefix α are irrelevant (recall that the i -th iteration reads exactly the i -th element of the input sequence α').

A key ingredient in the construction of the sequence α is the following lemma. Recall that for any set X and any integer $r \geq 1$, an r -coloring of X is simply a mapping $f : X \rightarrow [r]$.

Lemma 30 (Dominating color). *Let $r \geq 1$ be an integer and X_1, \dots, X_r finite, nonempty sets. For any r -coloring f of $X_1 \times \dots \times X_r$ there is a color $\sigma \in [r]$ such that for every $x_\sigma \in X_\sigma$ there are elements $x_s \in X_s$, $s \in [r] \setminus \{\sigma\}$, with $f(x_1, \dots, x_r) = \sigma$.*

We defer the proof of Lemma 30 to the next section.

Consider now the pseudocode description of ABSTRACTBUILD() in Algorithm 2. Note that its loop structure mirrors the structure of COMPUTEWEIGHTS(), with the crucial difference that while the loop (**) of COMPUTEWEIGHTS() simply focuses on one color $\sigma \in [r]$ (as indicated by the i -th entry of the input sequence α), for the strategy ABSTRACTBUILD() the 'right' color σ depends on the individual decisions of AbstractPainter occurring during the loop (++), and is therefore not known until this loop terminates.

In the next section we will prove the following two properties of ABSTRACTBUILD(F, r, θ).

Lemma 31 (Well-definedness and duration of AbstractBuilder strategy). *For F , r , and θ as in Proposition 6, the strategy ABSTRACTBUILD(F, r, θ) enforces a monochromatic copy of F in at most $r^2 \cdot |\mathcal{S}(F)|^{r+2}$ steps of the abstract game.*

Lemma 32 (AbstractBuilder strategy is legal). *For F , r , θ , and β as in Proposition 6, the strategy $\text{ABSTRACTBUILD}(F, r, \theta)$ satisfies the generalized density restriction (θ, β) .*

Together with Lemma 29, the preceding statements about $\text{ABSTRACTBUILD}(F, r, \theta)$ imply Proposition 6 straightforwardly.

Proof of Proposition 6. Using Lemma 31 and Lemma 32, we may apply Lemma 29 to $\text{ABSTRACTBUILD}(F, r, \theta)$ to obtain a strategy $\text{BUILD}(F, r, \theta)$ which enforces a monochromatic copy of F in the deterministic game with r colors in at most

$$a_{\max} = a_{\max}(F, r) := (r + 1)^{r^2 \cdot |S(F)|^{r+2}} \quad (97)$$

steps, and satisfies the generalized density restriction (θ, β) . \square

It remains to prove Lemma 30, Lemma 31, and Lemma 32, which we will do in the next section.

4.3. Analysis of $\text{ABSTRACTBUILD}()$.

Proof of Lemma 30. We refer to a color $\sigma \in [r]$ that satisfies the conditions of the lemma as a *color that is dominating in $X_1 \times \cdots \times X_r$* .

We argue by double induction over r and $\sum_{s \in [r]} |X_s|$. To settle the induction base note that the claim is trivially true for $r = 1$ and any finite set X_1 , and also for $r \geq 2$ and $|X_1| = \cdots = |X_r| = 1$. For the induction step let $r \geq 2$ and suppose that one of the sets X_s , $s \in [r]$, contains at least two elements. We assume w.l.o.g. that it is X_1 , and fix an element $x \in X_1$.

By induction (over the sum of the cardinalities of the sets X_s) we know that for the restriction of f to the set $(X_1 \setminus \{x\}) \times X_2 \times \cdots \times X_r$ there is a dominating color $\sigma \in [r]$. If $\sigma \neq 1$, then σ is also dominating in $X_1 \times \cdots \times X_r$ and we are done. Otherwise we have $\sigma = 1$, i.e. for all $x_1 \in X_1 \setminus \{x\}$ there are elements $x_s \in X_s$, $2 \leq s \leq r$, with $f(x_1, \dots, x_r) = 1$. Therefore, if f assigns color 1 to any of the elements in $\{x\} \times X_2 \times \cdots \times X_r$, then $\sigma = 1$ is dominating in $X_1 \times \cdots \times X_r$ and we are done as well. The only remaining case is that $f(x, \bullet, \dots, \bullet)$ never uses color 1, and therefore is an $(r - 1)$ -coloring of $X_2 \times \cdots \times X_r$. By induction (over r), there is a color $\sigma' \in [r] \setminus \{1\}$ that is dominating in $X_2 \times \cdots \times X_r$, and therefore also in $X_1 \times \cdots \times X_r$. This settles the last remaining case. \square

In order to prove Lemma 31 and Lemma 32 we will make use of the following technical lemma, which relates the evolution of the families \mathcal{G}_s occurring in $\text{ABSTRACTBUILD}()$ to the evolution of the families \mathcal{H}_s occurring in $\text{COMPUTEWEIGHTS}()$.

Lemma 33 (Evolution of the families \mathcal{G}_s). *At the end of each iteration of the repeat-loop $(++)$ during some iteration i of the repeat-loop $(+)$ in $\text{ABSTRACTBUILD}()$, for each $s \in [r]$ we have $\mathcal{G}_s \supseteq \mathcal{H}_s$, where \mathcal{H}_s denotes the value of \mathcal{H}_s after $j_s - 1$ iterations of the repeat-loop $(**)$ during iteration i of the repeat-loop $(*)$ in $\text{COMPUTEWEIGHTS}()$ for the input sequence $\alpha \circ s$, for the current value of j_s . Here $\alpha \in [r]^{i-1}$ denotes the sequence that has been constructed in previous rounds of $\text{ABSTRACTBUILD}()$ as a result of AbstractPainter's coloring decisions.*

Proof. Consider the i -th iteration of the repeat-loop $(+)$, and let $\hat{\sigma}$ be the color defined in line B14 in some iteration of the repeat-loop $(++)$. Note that the set of graphs (H, π) that are added to $\mathcal{G}_{\hat{\sigma}}$ in line B17 and line B23 in this iteration is

$$\mathcal{H}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}} := \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}} \cup \bigcup_{k \geq 1} \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}, k} \quad (98)$$

Algorithm 2: AbstractBuilder strategy $\text{ABSTRACTBUILD}(F, r, \theta)$

Input: a graph F with at least one edge, an integer $r \geq 2$, a real number $\theta > 0$

```

B1  $\alpha := ()$ 
B2 foreach  $s \in [r]$  do
B3    $\mathcal{G}_s := \emptyset$ 
B4  $i := 0$ 
B5 repeat (+)
B6    $i := i + 1$ 
B7   foreach  $s \in [r]$  do
B8     Let  $j_{\max, s}$  denote the total number of iterations of the repeat-loop (**) in the  $i$ -th
      iteration of the repeat-loop (*) of the algorithm  $\text{COMPUTEWEIGHTS}()$  with input
      sequence  $\alpha \circ s$ .
B9     For  $1 \leq j \leq j_{\max, s}$ , let  $\mathcal{C}_s^{i, j}$  and  $\mathcal{C}_s^{i, j, k}$ ,  $k \geq 1$ , denote the sets defined in the algorithm
       $\text{COMPUTEWEIGHTS}()$  with input sequence  $\alpha \circ s$  in the corresponding iterations of the
      repeat-loops (**) and (***) in line 14, or line 24 and 32, respectively.
B10     $j_s := 1$ 
B11  repeat (++)
B12    foreach  $((H_1, \pi_1), \dots, (H_r, \pi_r)) \in \mathcal{C}_1^{i, j_1} \times \dots \times \mathcal{C}_r^{i, j_r}$  do
B13      For each color  $s \in [r]$  let  $v_{s1}$  denote the youngest vertex of  $(H_s, \pi_s)$ . AbstractBuilder
      constructs a new graph by taking the disjoint union of all graphs
       $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ ,  $s \in [r]$ , for which  $v(H_s) \geq 2$ , and by adding a new vertex  $v$  in
      such a way that for each  $s \in [r]$ , coloring  $v$  in color  $s$  will extend the central copy of
       $H_s \setminus v_{s1}$  in  $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$  to a copy of  $H_s$ . AbstractPainter chooses a color
       $\bar{\sigma} \in [r]$  for  $v$ , the resulting new  $r$ -colored graph  $G$  is added to AbstractBuilder's list,
      and the newly created copy of  $H_{\bar{\sigma}}$  is designated as the central copy of  $G$ .
B14    By Lemma 30, for any combination of colors AbstractPainter chooses in the previous
      loop (line B12 and B13), there is a color  $\hat{\sigma} \in [r]$  such that for each graph  $(H, \pi) \in \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}}$ ,
      in at least one construction step an  $r$ -colored graph with a central copy of  $H$  in color  $\hat{\sigma}$ 
      is created. Fix such a color  $\hat{\sigma}$ .
B15    foreach  $(H, \pi) \in \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}}$  do
B16      Define  $G_{\hat{\sigma}}(H, \pi)$  to be the  $r$ -colored graph on AbstractBuilder's list resulting from
      an arbitrary construction step in line B13 that created a central copy of  $H$  in
      color  $\hat{\sigma}$ .
B17     $\mathcal{G}_{\hat{\sigma}} := \mathcal{G}_{\hat{\sigma}} \cup \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}}$ 
B18     $k := 0$ 
B19    repeat (+++)
B20       $k := k + 1$ 
B21      foreach  $(H, \pi) \in \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}, k}$ ,  $\pi = (v_1, \dots, v_h)$ , do
B22        By Lemma 23, the graph  $(H, \pi')$ , defined by  $\pi' := (v_{k+1}, v_1, \dots, v_k, v_{k+2}, \dots, v_h)$ 
        is contained in  $\mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}}$ , and consequently  $G_{\hat{\sigma}}(H, \pi')$  was just defined in line B16.
        Let  $G_{\hat{\sigma}}(H, \pi) := G_{\hat{\sigma}}(H, \pi')$ .
B23       $\mathcal{G}_{\hat{\sigma}} := \mathcal{G}_{\hat{\sigma}} \cup \mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}, k}$ 
B24    until  $\mathcal{C}_{\hat{\sigma}}^{i, j_{\hat{\sigma}}, k} = \emptyset$ 
B25     $j_{\hat{\sigma}} := j_{\hat{\sigma}} + 1$ 
B26  until  $j_{\sigma} > j_{\max, \sigma}$  for some  $\sigma \in [r]$ 
B27   $\alpha := \alpha \circ \sigma$  for this  $\sigma$ 
B28 until  $(F, \pi) \in \mathcal{G}_s$  for some  $s \in [r]$  and  $\pi \in \Pi(V(F))$ 

```

for the current value of $j_{\hat{\sigma}}$, where the union in (98) is over all $k \geq 1$ for which $\mathcal{C}_{\hat{\sigma}}^{i,j_{\hat{\sigma}},k}$ is defined in line B9. By the definitions in line B9 these are exactly the graphs that are added to the family $\mathcal{H}_{\hat{\sigma}}$ in the $j_{\hat{\sigma}}$ -th iteration of the repeat-loop (**) in the i -th iteration of the repeat-loop (*) of the algorithm COMPUTEWEIGHTS() for the input sequence $\alpha \circ \hat{\sigma}$. It follows inductively that at the end of the i -th iteration of the repeat-loop (+), the set of graphs (H, π) that have been added to \mathcal{G}_{σ} for the color $\sigma \in [r]$ satisfying the termination condition in line B26 is

$$\mathcal{H}_{\sigma}^i := \bigcup_{j_{\sigma}=1}^{j_{\max,\sigma}} \mathcal{H}_{\sigma}^{i,j_{\sigma}}.$$

By the definition of $j_{\max,s}$ in line B8, these are exactly the graphs added to the family \mathcal{H}_{σ} in the i -th iteration of the repeat-loop (*) of the algorithm COMPUTEWEIGHTS() for the input sequence $\alpha \circ \sigma$. As moreover no graphs are added to the families \mathcal{H}_s , $s \in [r] \setminus \{\sigma\}$, during the i -th round of COMPUTEWEIGHTS() with input sequence $\alpha \circ \sigma$, it follows inductively that throughout, the families \mathcal{G}_s occurring in ABSTRACTBUILD() are related to the families \mathcal{H}_s occurring in COMPUTEWEIGHTS() as claimed. \square

Proof of Lemma 31. We will first argue that ABSTRACTBUILD() is a well-defined winning strategy.

Note that whenever a graph $(H, \pi) \in \mathcal{S}(F)$ is added to one of the families \mathcal{G}_s , $s \in [r]$, in line B17 or B23, then $G_s(H, \pi)$ is defined in B16 or in line B22, respectively, and in either case this r -colored graph was added to AbstractBuilder's list in line B13. Thus throughout the strategy ABSTRACTBUILD(), for every $s \in [r]$ and all graphs $(H, \pi) \in \mathcal{G}_s$, the graph $G_s(H, \pi)$ is well-defined and exists on AbstractBuilder's list. With the termination condition in line B28 this implies in particular that when ABSTRACTBUILD() terminates, AbstractBuilder's list indeed contains a graph containing a monochromatic copy of F .

Next we show that whenever the construction step in line B13 is executed, all involved graphs $(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ are in the respective families \mathcal{G}_s at this point, and thus the graphs $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ used for the construction step are indeed on AbstractBuilder's list. It follows from the definition of $\mathcal{C}^{i,j}$ as a subset of $\mathcal{C}(\mathcal{H}_s, F)$ (recall (22)) in line 14 of COMPUTEWEIGHTS() that for each $s \in [r]$ and each $(H_s, \pi_s) \in \mathcal{C}_s^{i,j_s}$ with $v(H_s) \geq 2$ in line B12, the graph $(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ is contained in \mathcal{H}_s as defined in the i -th iteration of the repeat-loop (*) at the beginning of the j_s -th iteration of the repeat-loop (**) of COMPUTEWEIGHTS() for the input sequence $\alpha \circ s$. Thus by Lemma 33, the graph $(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ is in the corresponding family \mathcal{G}_s , as claimed.

Also note that whenever the construction step in line B13 is executed, the involved entries $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$, $s \in [r]$, on AbstractBuilder's list are different from each other, as the corresponding central copies of $H_s \setminus v_{s1}$, $s \in [r]$, are all in different colors.

Together the above arguments show that ABSTRACTBUILD() is indeed a well-defined strategy for the abstract game, and it remains to bound the number of construction steps ABSTRACTBUILD() needs to enforce a monochromatic copy of F . By Lemma 33 and the termination condition in line B28, the number of iterations of the repeat-loop (+) until ABSTRACTBUILD() terminates is bounded by the number of iterations of the repeat-loop (*) in COMPUTEWEIGHTS() until the first of the families \mathcal{H}_s , $s \in [r]$, contains the graph (F, π) for some vertex-ordering $\pi \in \Pi(V(F))$. The termination condition in line 36 and Lemma 12 therefore show that ABSTRACTBUILD() terminates after at most $r \cdot |\mathcal{S}(F)|$ iterations of the repeat-loop (+). In each iteration i , the number of iterations of the repeat-loop (++) is at most $\sum_{s \in [r]} j_{\max,s} \leq r \cdot |\mathcal{S}(F)|$, as the values $j_{\max,s}$ are bounded by $|\mathcal{S}(F)|$ as argued in the proof of Lemma 12 on page 25. Lastly, the number of iterations of the loop in line B12 is $|\mathcal{C}_1^{i,j_1}| \cdots |\mathcal{C}_r^{i,j_r}| \leq |\mathcal{S}(F)|^r$, as the sets \mathcal{C}_s^{i,j_s} are subsets of $\mathcal{S}(F)$. Multiplying those numbers yields the claimed bound on the total number of construction steps throughout the strategy. \square

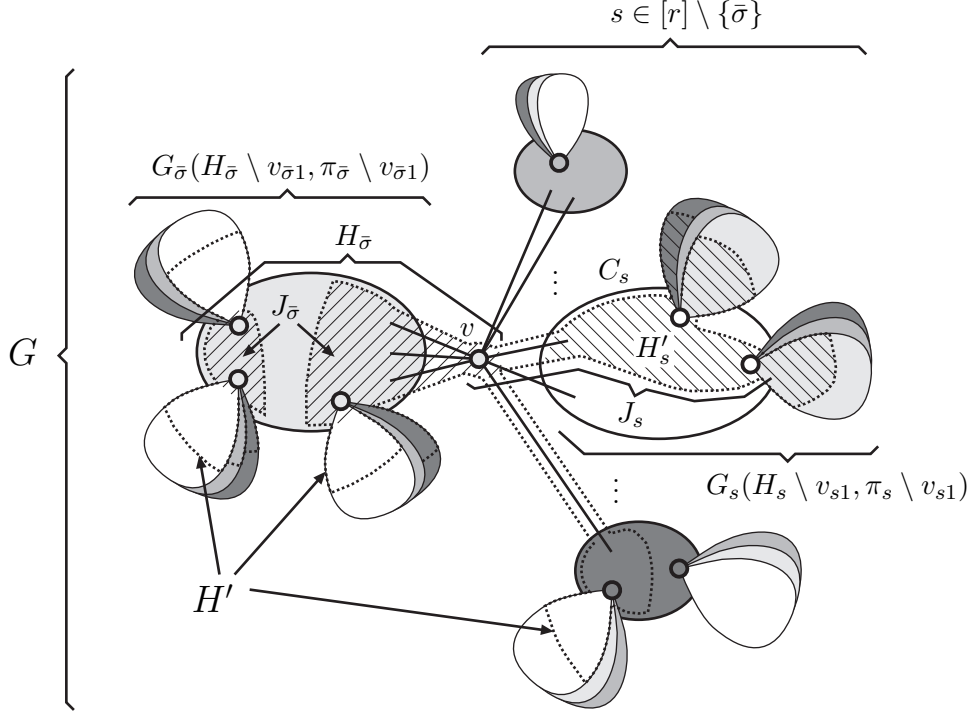


FIGURE 4. Notations used in the proof of Lemma 32.

Proof of Lemma 32. For the reader's convenience, Figure 4 illustrates the notations used throughout the proof.

We prove inductively that for each construction step in line B13 in some round i of **ABSTRACTBUILD**() the following holds: Recall the notations from the algorithm, and let H' , $v(H') \geq 1$, be a subgraph of the newly constructed graph G such that each component of H' shares at least one vertex with the central copy of $H_{\bar{\sigma}}$ in G . Letting $J_{\bar{\sigma}}$ denote the intersection of H' with this central copy, we have

$$\mu_{\theta}(H') \geq \lambda_{\theta}(J_{\bar{\sigma}}, w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}) \quad , \quad (99)$$

where here and throughout we denote for any $s \in [r]$ by $w_{(H, \pi, s), s}$ the weight function $w_{(H, \pi, s)}$ computed by **COMPUTEWEIGHTS**() for the input sequence $\alpha \circ s$. (Recall the definitions in (6), (19) and (23), and that during the i -th round of **ABSTRACTBUILD**(), the sequence α has length $i - 1$.)

For subgraphs $H' \subseteq G$ that do not contain the new vertex v , the claim follows by induction if $G_{\bar{\sigma}}(H_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}, \pi_{\bar{\sigma}} \setminus v_{\bar{\sigma}1})$ was defined in line B16 (either in the same or in an earlier iteration of the repeat-loop (+)), or by induction and by Lemma 23 if $G_{\bar{\sigma}}(H_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}, \pi_{\bar{\sigma}} \setminus v_{\bar{\sigma}1})$ was defined in line B22 (either in the same or in an earlier iteration of the repeat-loop (+)), recalling the definition (19) and the fact that the functions $w_{(H_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}, \pi_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}), \bar{\sigma}}$ and $w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}$ assign the same weight to all vertices of $H_{\bar{\sigma}}$ different from $v_{\bar{\sigma}1}$ (recall (23)).

It remains to prove (99) for subgraphs H' that contain the vertex v . Note that throughout the i -th iteration of the repeat-loop (+), whenever the construction step in line B13 is executed, by the definition of the sets \mathcal{C}_s^{i, j_s} (see line B9 of **ABSTRACTBUILD**() and line 14 of **COMPUTEWEIGHTS**()), all graphs $(H_s, \pi_s) \in \mathcal{C}_s^{i, j_s}$ used for the construction step satisfy

$$d_{\theta}(H_s, v_{s1}, w_{(H_s, \pi_s, s), s}) = d_s^i \quad , \quad s \in [r] \quad , \quad (100)$$

where the values d_s^i , $s \in [r]$, are defined in line 8 of COMPUTEWEIGHTS() with input sequence α . Note that these values depend only on the first $i - 1$ entries of α , i.e., they are the same during the i -th iteration of the repeat-loop (*) of COMPUTEWEIGHTS() for each input sequence $\alpha \circ s$, $s \in [r]$.

For the weight assigned to the youngest vertex $v_{\bar{\sigma}1}$ of $(H_{\bar{\sigma}}, \pi_{\bar{\sigma}})$ by COMPUTEWEIGHTS() with input sequence $\alpha \circ \bar{\sigma}$, we thus obtain by combining (23) with the definitions in line 10 and 18 that

$$w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}(v_{\bar{\sigma}1}) = \sum_{s \in [r] \setminus \{\bar{\sigma}\}} d_s^{i(100)} = \sum_{s \in [r] \setminus \{\bar{\sigma}\}} d_{\theta}(H_s, v_{s1}, w_{(H_s, \pi_s, s), s}) . \quad (101)$$

For each $s \in [r]$ with $v(H_s) \geq 2$ we define the graph H'_s as the intersection of H' with the copy of $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ used for the construction of G . Furthermore, for each such $s \in [r]$ we define a subgraph $J_s \subseteq H_s$ with $v_{s1} \in J_s$ as follows: Let C_s denote the central copy of $H_s \setminus v_{s1}$ in the copy of $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ used for the construction of G , and recall that the new vertex v completes C_s to a copy of H_s . Let $J_s \subseteq H_s$ denote the graph that is isomorphic to the intersection of H' with this copy of H_s , and note that H'_s intersects C_s in a copy of $J_s \setminus v_{s1}$. For all $s \in [r]$ with $v(H_s) = 1$ (i.e., H_s consists only of an isolated vertex) we define H'_s as the null graph (the graph whose vertex set is empty) and set $J_s := H_s$. Using these definitions we obtain

$$v(H') = \sum_{s \in [r]} v(H'_s) + 1 , \quad (102a)$$

$$e(H') = \sum_{s \in [r]} (e(H'_s) + \deg_{J_s}(v_{s1})) . \quad (102b)$$

Furthermore, for every $s \in [r]$ we have

$$\mu_{\theta}(H'_s) \geq \lambda_{\theta}(J_s \setminus v_{s1}, w_{(H_s, \pi_s, s), s}) \quad (103)$$

(this holds trivially if $v(H_s) = 1$; otherwise, similarly to before, if $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ was defined in line B16 then this follows by induction, whereas if $G_s(H_s \setminus v_{s1}, \pi_s \setminus v_{s1})$ was defined in line B22 then this follows by induction and by Lemma 23).

Combining our previous observations we obtain

$$\begin{aligned} \mu_{\theta}(H') &\stackrel{(6), (102)}{=} \sum_{s \in [r]} (\mu_{\theta}(H'_s) - \deg_{J_s}(v_{s1}) \cdot \theta) + 1 \\ &\stackrel{(103)}{\geq} \lambda_{\theta}(J_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}, w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}) - \deg_{J_{\bar{\sigma}}}(v_{\bar{\sigma}1}) \cdot \theta \\ &\quad + \sum_{s \in [r] \setminus \{\bar{\sigma}\}} \underbrace{(\lambda_{\theta}(J_s \setminus v_{s1}, w_{(H_s, \pi_s, s), s}) - \deg_{J_s}(v_{s1}) \cdot \theta)}_{\stackrel{(18), (19)}{\geq} d_{\theta}(H_s, v_{s1}, w_{(H_s, \pi_s, s), s})} + 1 \\ &\stackrel{(101)}{\geq} \lambda_{\theta}(J_{\bar{\sigma}} \setminus v_{\bar{\sigma}1}, w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}) - \deg_{J_{\bar{\sigma}}}(v_{\bar{\sigma}1}) \cdot \theta + w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}(v_{\bar{\sigma}1}) + 1 \\ &\stackrel{(20)}{=} \lambda_{\theta}(J_{\bar{\sigma}}, w_{(H_{\bar{\sigma}}, \pi_{\bar{\sigma}}, \bar{\sigma}), \bar{\sigma}}) , \end{aligned}$$

completing the inductive proof of (99).

From (99) it follows in particular that for every graph G that is added to AbstractBuilder's list during the i -th iteration of the repeat-loop (+), every connected subgraph $H' \subseteq G$ containing the last added vertex v satisfies

$$\mu_{\theta}(H') \geq \min_{J \subseteq H: v_1 \in J} \lambda_{\theta}(J, w_{(H, \pi, s), s}) \quad (104)$$

for some $s \in [r]$ and some (H, π) , $\pi = (v_1, \dots, v_h)$, from one of the sets $\mathcal{C}^{i,j}$ defined in the i -th iteration of COMPUTEWEIGHTS() when called with input sequence $\alpha \circ s$.

As argued in the proof of Lemma 28 (see (90a)), the right hand side of (104) equals $1 + \sum_{s \in [r]} d_s^i$ for the values d_s^i defined by COMPUTEWEIGHTS() with input sequence α . Regardless of how the sequence α constructed by ABSTRACTBUILD() evolves in further iterations of the repeat-loop (+), this quantity is decreasing in i by the first part of Lemma 16. Moreover, by Lemma 33 and the termination condition in line B28, ABSTRACTBUILD() terminates after at most \tilde{i} iterations, for \tilde{i} as defined in Lemma 28. It follows that all connected subgraphs H' with $v(H') \geq 1$ of all graphs G added to AbstractBuilder's list in the course of ABSTRACTBUILD() satisfy

$$\mu_\theta(H') \geq 1 + \sum_{s \in [r]} d_s^{\tilde{i}} , \quad (105)$$

where \tilde{i} and $d_s^{\tilde{i}}$, $s \in [r]$, are defined in COMPUTEWEIGHTS() for the input sequence α constructed by ABSTRACTBUILD().

By Lemma 28 and the definition of $\Lambda_\theta()$ in (24), we thus obtain from (105) that

$$\mu_\theta(H') \geq \Lambda_\theta(F, r) \stackrel{(9)}{\geq} \beta$$

for all connected subgraphs H' with $v(H') \geq 1$ of all graphs G appearing on AbstractBuilder's list. Due to the assumption that $\beta \geq 0$, the same statement also follows for all disconnected subgraphs $H' \subseteq G$ with $v(H') \geq 1$, concluding the proof that the strategy ABSTRACTBUILD(F, r, θ) respects the generalized density restriction (θ, β) throughout. \square

5. PAINTER IN THE DETERMINISTIC GAME

In this section we prove Proposition 7 by explicitly constructing, for F, r, θ and β as in the proposition, a Painter strategy that avoids creating a monochromatic copy of F in the deterministic game with r colors and generalized density restriction (θ, β) .

5.1. Painter's strategy and proof of Proposition 7. Consider the following Painter strategy, which has four parameters: a graph F with at least one edge, an integer $r \geq 2$, a real number $\theta > 0$ and a sequence $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$. The strategy uses the output of Algorithm 1: $((\mathcal{H}_s, w_s))_{s \in [r]} := \text{COMPUTEWEIGHTS}(F, r, \theta, \alpha)$. In each step of the game, Painter picks a color as follows: Let v denote the vertex added in the current step, and for each $s \in [r]$, define

$$\mathcal{D}_s := \left\{ (H, \pi) \in \mathcal{S}(F) \mid \begin{array}{l} \text{assigning color } s \text{ to } v \text{ would create a copy of } \\ (H, \pi) \text{ in color } s \text{ on the board} \end{array} \right\} . \quad (106)$$

(Note that this requires Painter to memorize the order in which the vertices on the board arrived.) Calculate for each color $s \in [r]$ the value

$$d(s) := \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)}) , \quad (107)$$

where $\lambda_\theta()$ is defined in (19), and $w_{(H, \pi, s)}()$ is defined in (23) using $\mathcal{H}_s \subseteq \mathcal{S}(F)$ and $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$ as returned by Algorithm 1. (It is possible that $d(s) = -\infty$ for some colors $s \in [r]$.) Then select $\sigma \in [r]$ as the color for which this value is maximal, and assign color σ to the vertex v .

Intuitively, the parameter $\lambda_\theta(H, w_{(H, \pi, s)})$ measures the 'level of danger' that the Painter strategy encoded by α assigns to copies of the ordered graph (H, π) in color s , where a graph is considered the more dangerous the smaller its $\lambda_\theta()$ -value is. Thus the definition of $d(s)$ in (107) corresponds to determining the most dangerous graph in color s that would be created by assigning color s to v , and our strategy selects σ as the color for which this most dangerous graph is least dangerous.

If several colors have the same maximal value of $d(s)$, the above rule does not determine a color σ uniquely. Such ties are broken as follows: Consider the families

$$\mathcal{H}'_s = \mathcal{H}'_s(F, r, \theta, \alpha) := \bigcup_{\alpha_i = s \wedge (i=1 \vee \alpha_i \neq \alpha_{i-1})} \mathcal{C}^{i,1}, \quad (108)$$

$s \in r$, where $\mathcal{C}^{i,j}$ are the sets defined in line 14 of the algorithm `COMPUTEWEIGHTS`(F, r, θ, α). Note that these families are fixed throughout Painter's strategy. In Lemma 50 below we will show that ties can arise only between *two* different colors, and that whenever such a tie arises, then for exactly one of the two colors the set

$$\mathcal{J}_s := \arg \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)}) \quad (109)$$

contains an ordered graph from the corresponding family \mathcal{H}'_s . Our tie-breaking rule is to then pick the *other* color, i.e., the color $\sigma \in [r]$ for which \mathcal{J}_σ contains no graph from \mathcal{H}'_σ . (Intuitively, Painter considers the ordered graphs in the families \mathcal{H}'_s , $s \in [r]$, as slightly more dangerous than other ordered graphs with the same $\lambda_\theta()$ -value.)

In the following we denote the Painter strategy defined above by `PAINT`(F, r, θ, α). Note that this strategy can be employed both in the deterministic two-player game and in the original probabilistic process.

Remark 34. Note that the actual $\lambda_\theta()$ -values of monochromatic ordered subgraphs of F are not relevant in the above strategy — all that matters is the partial order on the set $\mathcal{S}(F) \times [r]$ induced by the $\lambda_\theta()$ -values and our tie-breaking rule. This partial order can be extended arbitrarily to a total order by defining an arbitrary order among all elements of $\mathcal{S}(F) \times [r]$ that have the same $\lambda_\theta()$ -value and are in one of the sets \mathcal{H}'_s , and among all elements that have the same $\lambda_\theta()$ -value and are *not* in one of the sets \mathcal{H}'_s . Thus the strategy `PAINT`(F, r, θ, α) can indeed be represented as a priority list of ordered monochromatic subgraphs of F , as described in Section 1.5.

A careful analysis of the strategy `PAINT`(F, r, θ, α) will eventually yield the following key lemma. As its statement is purely deterministic, it is applicable to both the deterministic game and the probabilistic process. Note that the lemma does not assume any density restrictions for the evolving board.

Lemma 35 (Witness graph invariant). *For F , r , θ , and α as specified in Algorithm 1 there is a constant $v_{\max} = v_{\max}(F, r, \theta, \alpha)$ such that if Painter plays according to the strategy `PAINT`(F, r, θ, α) then the following invariant is maintained throughout:*

The board contains a graph K' with $v(K') \leq v_{\max}$ and

$$\mu_\theta(K') < 0, \quad (110)$$

or for every $s \in [r]$ and every $(H, \pi) \in \mathcal{S}(F)$ we have that every copy of (H, π) in color s on the board is contained in a graph H' with $v(H') \leq v_{\max}$ and

$$\mu_\theta(H') \leq \lambda_\theta(H, w_{(H, \pi, s)}) , \quad (111)$$

where $\mu_\theta()$, $\lambda_\theta()$, and $w_{(H, \pi, s)}$ are defined in (6), (19), and (23) (using $\mathcal{H}_s \subseteq \mathcal{S}(F)$ and $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$ as returned by Algorithm 1), respectively.

Remark 36. As we shall see shortly, the statement that the size of the graphs K' and H' in Lemma 35 is bounded by some constant $v_{\max} = v_{\max}(F, r, \theta, \alpha)$ is not needed to prove Proposition 7. However, it will be crucial for proving the lower bound part of Theorem 4 in Section 6.1 below (recall the remarks in Section 2.2 and Section 2.3). In fact, the proof of the existence of the bound v_{\max} relies primarily on our tie-breaking rule described above; a version of Lemma 35 *without* a bound on the size of the graphs K' and H' (which suffices to infer Proposition 7) can also be proven if ties are broken arbitrarily. In this case the proof of Lemma 35 can be simplified considerably; in particular,

Lemma 47, Lemma 48, Lemma 50 and the second part of Lemma 51 below are not needed. The reader might want to skip those parts on his first read-through, or if he is only interested in the deterministic game.

With Lemma 35 in hand, the proof of Proposition 7 is straightforward.

Proof of Proposition 7. Let $\alpha \in [r]^{r \cdot |\mathcal{S}(F)|}$ be a sequence for which the minimum in the definition of $\Lambda_\theta(F, r)$ in (24) is attained. By the definition in (24), for all colors $s \in [r]$ and all vertex orderings $\pi \in \Pi(V(F))$ there is a subgraph $H \subseteq F$ with

$$\lambda_\theta(H, w_{(H, \pi|_H, s)}) \leq \Lambda_\theta(F, r) \stackrel{(10)}{<} \beta . \quad (112)$$

Suppose now that Painter plays according to the strategy $\text{PAINT}(F, r, \theta, \alpha)$ and that for some $\pi \in \Pi(V(F))$ a copy of (F, π) in some color $s \in [r]$ appears on the board. Choose $H \subseteq F$ such that (112) holds. Then, by Lemma 35, the board contains a graph K' with

$$\mu_\theta(K') \stackrel{(110)}{<} 0 \leq \beta ,$$

or the copy of $(H, \pi|_H)$ in color s that is contained in the copy of (F, π) is contained in a graph H' with

$$\mu_\theta(H') \stackrel{(111)}{\leq} \lambda_\theta(H, w_{(H, \pi|_H, s)}) \stackrel{(112)}{<} \beta .$$

None of the two cases can occur if Builder adheres to the generalized density restriction (θ, β) , and consequently Painter can avoid creating a monochromatic copy of F in the deterministic F -avoidance game with r colors and generalized density restriction (θ, β) by playing according to the strategy $\text{PAINT}(F, r, \theta, \alpha)$. \square

The rest of this section is devoted to proving Lemma 35. To do so we will need a number of technical lemmas. Throughout the following, F , r , θ and α are fixed, and we usually omit these arguments when we refer to $\text{COMPUTEWEIGHTS}(F, r, \theta, \alpha)$ or $\text{PAINT}(F, r, \theta, \alpha)$. We let $\mathcal{H}_s \subseteq \mathcal{S}(F)$ and $w_s : \mathcal{H}_s \rightarrow \mathbb{R}$ denote the return values of $\text{COMPUTEWEIGHTS}()$, and $w_{(H, \pi, s)}$ the weight function defined in (23) with respect to these return values.

5.2. A geometric viewpoint. We begin by relating the strategy $\text{PAINT}()$ and many of the quantities defined in previous parts of this paper to a simple geometric object. This geometric viewpoint will be a key ingredient in our proof of Lemma 35.

Definition 37 (Axis-parallel decreasing walk). We say that $(x_\nu)_{1 \leq \nu \leq k}$, $x_\nu \in \mathbb{R}^r$, is a *decreasing axis-parallel walk in \mathbb{R}^r* if for any two subsequent elements x_ν and $x_{\nu+1}$ there is a coordinate $s \in [r]$ such that $x_{\nu+1, s} < x_{\nu, s}$ and $x_{\nu+1, t} = x_{\nu, t}$ for all $t \in [r] \setminus \{s\}$.

The following lemma is an immediate consequence of this definition.

Lemma 38 (Order on the walk). *Let $(x_\nu)_{1 \leq \nu \leq k}$, $x_\nu \in \mathbb{R}^r$, be a decreasing axis-parallel walk in \mathbb{R}^r . For any two elements x_μ, x_ν we have $1 + \sum_{t \in [r]} x_{\nu, t} \leq 1 + \sum_{t \in [r]} x_{\mu, t}$ if and only if $x_{\nu, t} \leq x_{\mu, t}$ for all $t \in [r]$.*

We can think of the points $(x_\nu)_{1 \leq \nu \leq k}$ of a decreasing axis-parallel walk as lying on a sequence of consecutive axis-parallel line segments. For technical reasons we also specify a direction $\sigma \in [r]$ in which, intuitively, the walk continues beyond the point x_k . Moreover, we sometimes allow the last segment of this walk to degenerate into the single point x_k . This is made precise in the following definition.

Definition 39 (Extended walk, turning point, starting point/endpoint, segment, order). Given some decreasing axis-parallel walk $(x_\nu)_{1 \leq \nu \leq k}$, $x_\nu \in \mathbb{R}^r$, and some $\sigma \in [r]$, we say that the pair $((x_\nu)_{1 \leq \nu \leq k}, \sigma)$ is an *extended decreasing axis-parallel walk in \mathbb{R}^r* . We refer to a point x_ν , $2 \leq \nu \leq k-1$, on this extended walk as a *turning point* if the coordinate in which x_ν differs from $x_{\nu-1}$ is different from the coordinate in which x_ν differs from $x_{\nu+1}$. We also define x_1 to be a turning point. Moreover, x_k is a turning point if and only if the coordinate in which x_k differs from x_{k-1} is different from σ .

Given two consecutive turning points x_μ and x_ν , $\mu < \nu$, that differ in some coordinate $s \in [r]$, we refer to the line segment that connects them as an *s-segment*, and we call x_μ the *starting point* and x_ν the *endpoint* of this segment. Furthermore, we refer to the line segment that connects the last turning point x_μ to the last point x_k of the walk as a σ -segment (if $\mu = k$, this segment degenerates to a single point x_k). We call x_μ the starting point and x_k the endpoint of this segment.

For two points x_μ and x_ν , we say that x_μ is *higher on the walk than x_ν* (or equivalently, x_ν is lower on the walk than x_μ) if $\mu < \nu$. We extend this notion to segments on the extended walk by saying that an *s-segment* Γ is higher on the walk than an *s'-segment* Γ' (or equivalently, Γ' is lower than Γ) if the starting point of Γ is higher on the walk than the starting point of Γ' .

Note that according to Lemma 16, the points (d_1^i, \dots, d_r^i) , $i \geq 1$, form a decreasing axis-parallel walk in \mathbb{R}^r . In the following we define, for every $s \in [r]$ and every $(H, \pi) \in \mathcal{H}_s$, a point $x_{(H, \pi, s)} \in \mathbb{R}^r$ on one of the line segments of this walk.

Let d_s^i , $s \in [r]$, denote the values defined in line 8 of the algorithm COMPUTEWEIGHTS(), and $\mathcal{C}^{i,j}$ and $\mathcal{C}^{i,j,k}$ the sets defined in line 14, or line 24 and 32, respectively. Recall from Section 3.3 that for each $s \in [r]$, the sets $\mathcal{C}^{i,j}$ and $\mathcal{C}^{i,j,k}$ for which $\alpha_i = s$ form a partition of the family \mathcal{H}_s .

Definition 40 (*x*-points). For every $s \in [r]$ and every graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, we define a point $x_{(H, \pi, s)} \in \mathbb{R}^r$ as follows:

(a) If $(H, \pi) \in \mathcal{C}^{i,j}$ for some $i, j \geq 1$ with $\alpha_i = s$, then we define

$$x_{(H, \pi, s)} := (d_1^i, \dots, d_{s-1}^i, d_\theta(H, v_1, w_{(H, \pi, s)}) = d_s^i, d_{s+1}^i, \dots, d_r^i) \in \mathbb{R}^r. \quad (113a)$$

(b) If $(H, \pi) \in \mathcal{C}^{i,j,k}$ for some $i, j, k \geq 1$ with $\alpha_i = s$, then we define

$$x_{(H, \pi, s)} := (d_1^{\hat{i}}, \dots, d_{s-1}^{\hat{i}}, d_\theta(H, v_1, w_{(H, \pi, s)}) = d_s^{\hat{i}}, d_{s+1}^{\hat{i}}, \dots, d_r^{\hat{i}}) \in \mathbb{R}^r, \quad (113b)$$

where \hat{i} is the value defined either in line 28 or in line 30 for this (H, π) .

The fact that $d_\theta(H, v_1, w_{(H, \pi, s)}) = d_s^i$ in part (a) of the previous definition follows directly from the definition in line 14 of the algorithm COMPUTEWEIGHTS().

The following lemma states that the points $x_{(H, \pi, s)}$ defined above indeed form an (extended) decreasing axis-parallel walk. We point out that the order in which the points $x_{(H, \pi, s)}$ appear on this walk is not necessarily the order in which the corresponding graphs (H, π) are added to the families \mathcal{H}_s in the course of the algorithm COMPUTEWEIGHTS(). To be more precise, such a statement is true for the graphs that are added via one of the sets $\mathcal{C}^{i,j}$, but not for the graphs that are added via one of the sets $\mathcal{C}^{i,j,k}$.

Of particular importance are the points $x_{(H, \pi, s)}$ for the graphs $(H, \pi) \in \mathcal{H}'_s$, where $\mathcal{H}'_s \subseteq \mathcal{H}_s$ are the families defined in (108) for our tie-breaking rule. As it turns out, those points are always turning points of the walk. Figure 5 illustrates Definition 40 and the different statements of Lemma 41.

Lemma 41 (Walk formed by *x*-points). *Let i_{\max} denote the total number of iterations of the repeat-loop (*) of COMPUTEWEIGHTS(). The elements in the set $\{x_{(H, \pi, s)} \mid s \in [r] \wedge (H, \pi) \in \mathcal{H}_s\}$ can be ordered to form a decreasing axis-parallel walk $W = W(F, r, \theta, \alpha)$ in \mathbb{R}^r such that the extended walk $\mathcal{W} = \mathcal{W}(F, r, \theta, \alpha) := (W, \alpha_{i_{\max}})$ satisfies the following properties:*

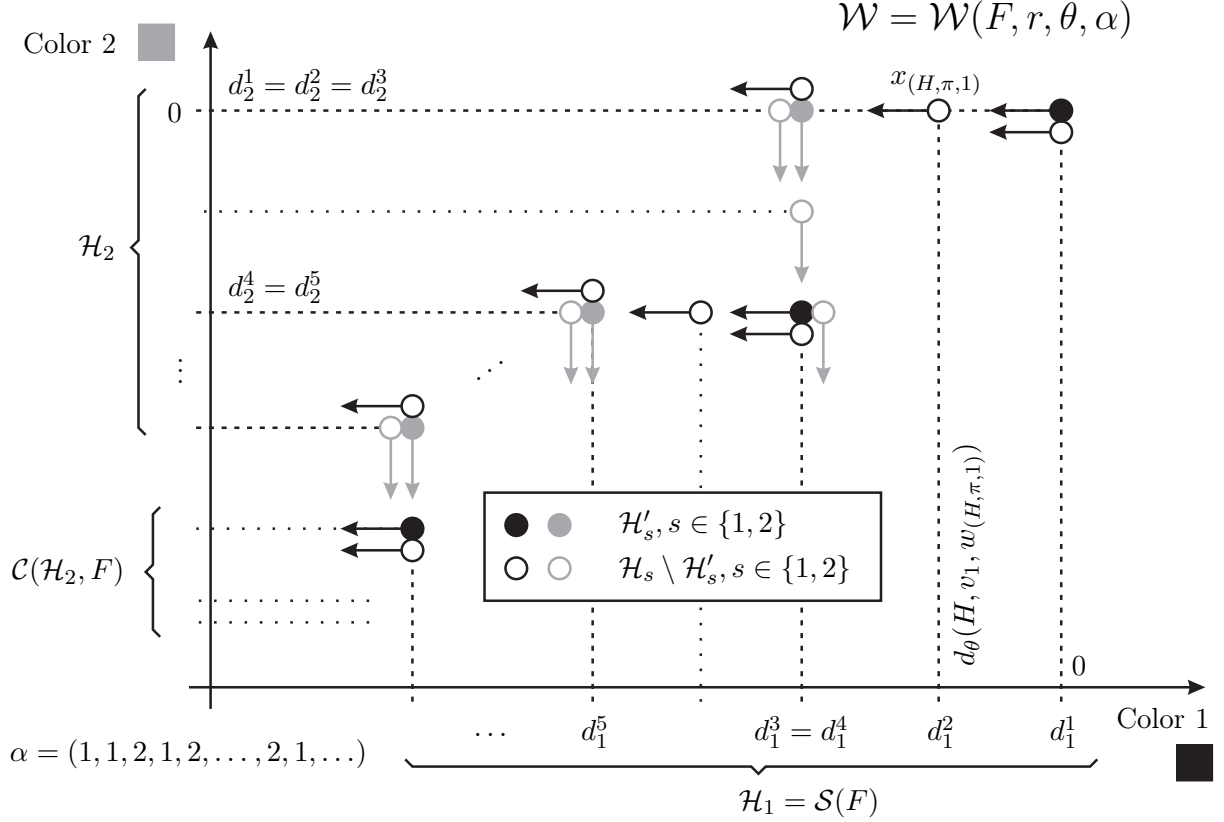


FIGURE 5. Illustration of Definition 40 and Lemma 41. The figure shows how certain variables of the algorithm $\text{COMPUTEWEIGHTS}(F, r, \theta, \alpha)$ might evolve for some graph F , some value of θ , $r = 2$ and $\alpha = (1, 1, 2, 1, 2, \dots, 2, 1, \dots)$ (where the algorithm terminates in the round corresponding to the last 1-entry shown). For any graph $(H, \pi) \in \mathcal{H}_s$, $s \in \{1, 2\}$, a bullet shows the location of the point $x_{(H, \pi, s)}$, and the arrow attached to the bullet points along the s -axis. One bullet may represent multiple points at the same location if the corresponding graphs are of the same type (see legend). If graphs with the same associated point are of different type, then the bullets are drawn directly adjacent to each other (instead of on top of each other) to maintain readability.

- (i) For any $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$, the point $x_{(H, \pi, s)}$ is contained in an s -segment.
- (ii) For any s -segment Γ , there is a graph $(H, \pi) \in \mathcal{H}_s$ such that $x_{(H, \pi, s)}$ is the starting point of Γ .
- (iii) For any s -segment Γ , if $x_{(H, \pi, s)}$, $\pi = (v_1, \dots, v_h)$, is the starting point of Γ , then there is some $J \subseteq H$ with $v_1 \in J$ such that $x_{(J, \pi|_J, s)} = x_{(H, \pi, s)}$ and $(J, \pi|_J) \in \mathcal{H}'_s$.
- (iv) For any s -segment Γ , if $x_{(H, \pi, s)} \in \Gamma$ is not the starting point of Γ , then $(H, \pi) \notin \mathcal{H}'_s$.
- (v) The lowest segment of \mathcal{W} is an $\alpha_{i_{\max}}$ -segment and $\mathcal{H}_{\alpha_{i_{\max}}} = \mathcal{S}(F)$.

Proof. By the first part of Lemma 16, the sequence $\overline{W} := (\overline{W}^i)_{1 \leq i \leq i_{\max}}$, $\overline{W}^i := (d_1^i, \dots, d_r^i)$, is a decreasing axis-parallel walk with the property that \overline{W}^i and \overline{W}^{i+1} differ exactly in the coordinate α_i . We first show that the set

$$\{x_{(H, \pi, s)} \mid s \in [r] \wedge (H, \pi) \in \mathcal{C}^{i, j} \text{ for some } i, j \geq 1 \text{ with } \alpha_i = s\} \quad (114)$$

coincides exactly with the set of elements of this walk, and that the extended walk $\overline{W} := (\overline{W}, \alpha_{i_{\max}})$ satisfies the properties of the lemma. In the second part of the proof we argue that for any $s \in [r]$ and any graph $(H, \pi) \in \mathcal{C}^{i,j,k}$ for some $i, j, k \geq 1$ with $\alpha_i = s$, the point $x_{(H, \pi, s)}$ lies on one of the segments of \overline{W} , and that the subdivided walk obtained by inserting all those points into the walk \overline{W} still satisfies the claimed properties.

First note that on the extended walk \overline{W} , every point \overline{W}^i , $1 \leq i \leq i_{\max}$, is contained in an α_i -segment. For any $s \in [r]$, any $i, j \geq 1$ with $\alpha_i = s$ and any graph $(H, \pi) \in \mathcal{C}^{i,j}$, by the definition in (113a) we have

$$x_{(H, \pi, s)} = (d_1^i, \dots, d_r^i) = \overline{W}^i \quad (115)$$

(independently of j). Thus property (i) is satisfied for the elements in the set (114) and the walk \overline{W} .

Recall that for each $1 \leq i \leq i_{\max}$ the set $\mathcal{C}^{i,1}$ is nonempty (see the definitions in line 8 and line 14), implying that there is a graph $(H, \pi) \in \mathcal{C}^{i,1}$ which for $s := \alpha_i$ satisfies $x_{(H, \pi, s)} = \overline{W}^i$, proving in particular property (ii) for the walk \overline{W} .

To prove properties (iii) and (iv), we fix some $s \in [r]$, some $i, j \geq 1$ with $\alpha_i = s$ and some graph $(H, \pi) \in \mathcal{C}^{i,j}$, $\pi = (v_1, \dots, v_h)$. Let Γ denote the s -segment of \overline{W} containing the point $x_{(H, \pi, s)}$. We distinguish two cases depending on whether $x_{(H, \pi, s)}$ is the starting point of Γ or not. Note that by the first part of Lemma 16, $x_{(H, \pi, s)} = \overline{W}^i$ is the starting point of Γ if and only if $i = 1$ or $\alpha_i \neq \alpha_{i-1}$.

We first consider the case that $x_{(H, \pi, s)}$ is the starting point of Γ , i.e., we have

$$i = 1 \text{ or } \alpha_i \neq \alpha_{i-1} . \quad (116)$$

If $j = 1$, then by (108) and (116) we have $(H, \pi) \in \mathcal{H}'_s$. If $j > 1$, then by the second part of Lemma 17 (recall that by the definition in line 14 we have $d_\theta(H, v_1, w_{(H, \pi, s)}) = d_s^i$) there is a subgraph $J \subseteq H$ with $v_1 \in J$ for which $(J, \pi|_J)$ is contained in $\mathcal{C}_s(d_s^i)$. By the definition in line 16 we have $\mathcal{C}_s(d_s^i) = \mathcal{C}^{i,1}$, implying that $(J, \pi|_J) \in \mathcal{C}^{i,1}$ and

$$x_{(J, \pi|_J, s)} \stackrel{(115)}{=} x_{(H, \pi, s)} .$$

Furthermore, using (116) it follows from the definition in (108) that $(J, \pi|_J) \in \mathcal{H}'_s$, proving that property (iii) holds for the walk \overline{W} .

If on the other hand $x_{(H, \pi, s)}$ is not the starting point of Γ , i.e., $i > 1$ and $\alpha_i = \alpha_{i-1}$, then by the definition in (108) we have $(H, \pi) \notin \mathcal{H}'_s$, proving property (iv) for the walk \overline{W} .

Note that by the definition of \overline{W} , the lowest segment of this walk is indeed an $\alpha_{i_{\max}}$ -segment. By the termination condition in line 36 and the observation that during the i -th iteration of the repeat-loop (*), none of the families \mathcal{H}_s , $s \in [r] \setminus \{\alpha_i\}$, is modified, we have $\mathcal{H}_{\alpha_{i_{\max}}} = \mathcal{S}(F)$. Together this proves property (v) for the walk \overline{W} .

To complete the proof of the lemma we fix some $s \in [r]$, some $i, j, k \geq 1$ with $\alpha_i = s$ and some graph $(H, \pi) \in \mathcal{C}^{i,j,k}$, $\pi = (v_1, \dots, v_h)$, and show that the point $x_{(H, \pi, s)}$ lies on some s -segment of the walk \overline{W} (possibly in between two points \overline{W}^i and \overline{W}^{i+1}), and that by including all such points $x_{(H, \pi, s)}$ into the walk \overline{W} we obtain a subdivided walk \mathcal{W} that still satisfies the claimed properties (note that beside (i) we only need to verify that properties (iii) and (iv) are maintained).

Note that by the definitions in line 28 and line 30 we have $\alpha_{\hat{i}} = \alpha_i = s$ for \hat{i} as in part (b) of Definition 40. Using this relation, the definition in (113b), and Lemma 19 we obtain that $x_{(H, \pi, s)}$ lies on the s -segment Γ that contains $\overline{W}^{\hat{i}}$ and $\overline{W}^{\hat{i}+1}$ on the walk \overline{W} , showing that the walk \mathcal{W} satisfies property (i).

By the strict inequality in (26a), $x_{(H,\pi,s)}$ can not be the starting point of Γ if \hat{i} was defined in line 28. Moreover, by (26b) the point $x_{(H,\pi,s)}$ is the starting point of Γ if and only if

$$d_\theta(H, v_1, w_{(H,\pi,s)}) = d_s^{\hat{i}} \quad (117)$$

and

$$\hat{i} = 1 \text{ or } \alpha_i \neq \alpha_{i-1} . \quad (118)$$

In this case, by the condition in line 27 there is a subgraph $J \subseteq H$ with $v_1 \in J$ such that $(J, \pi|_J)$ is contained in $\mathcal{C}_s(d_\theta(H, v_1, w_{(H,\pi,s)}))$. Using (117) and the definition in line 16 shows that $(J, \pi|_J) \in \mathcal{C}^{\hat{i},1}$, implying that

$$x_{(J,\pi|_J,s)} \stackrel{(115)}{=} (d_1^{\hat{i}}, \dots, d_r^{\hat{i}}) \stackrel{(113b),(117)}{=} x_{(H,\pi,s)} .$$

Furthermore, using (118) it follows from the definition in (108) that $(J, \pi|_J) \in \mathcal{H}'_s$, proving that property (iii) holds for the walk \mathcal{W} .

As none of the graphs in the sets $\mathcal{C}^{i,j,k}$ with $\alpha_i = s$ is contained in \mathcal{H}'_s (recall (108)), the walk \mathcal{W} trivially satisfies property (iv). This completes the proof. \square

5.3. Relation of the walk to other quantities. In the following lemmas we establish several relations between the walk \mathcal{W} defined in Lemma 41, the parameters $d_\theta()$ and $w_{(H,\pi,s)}$ used in the algorithm COMPUTEWEIGHTS(), and the parameter $\lambda_\theta()$ and the families \mathcal{H}'_s used in the definition of the strategy PAINT(). We will see that for the ordered monochromatic subgraphs of F that are relevant for the strategy PAINT(), the order of the corresponding x -points on the walk \mathcal{W} coincides with the ordering given by the $\lambda_\theta()$ -values — the lower on the walk the point $x_{(H,\pi,s)}$ appears, the lower the value $\lambda(H, w_{(H,\pi,s)})$, i.e., the more dangerous a copy of (H, π) in color s is considered (see Lemma 45 below).

Lemma 42 ($d_\theta()$ -value and weight from x -point). *For any $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, we have*

$$d_\theta(H, v_1, w_{(H,\pi,s)}) = x_{(H,\pi,s),s}$$

and

$$w_{(H,\pi,s)}(v_1) = \sum_{t \in [r] \setminus \{s\}} x_{(H,\pi,s),t} . \quad (119)$$

Proof. The first part of the lemma is an immediate consequence of the definition in (113).

For any $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$ as in part (a) of Definition 40 we obtain, using the definitions in line 10 and line 18,

$$w_{(H,\pi,s)}(v_1) \stackrel{(23)}{=} w_s(H, \pi) = w^i = \sum_{t \in [r] \setminus \{s\}} d_t^i \stackrel{(113a)}{=} \sum_{t \in [r] \setminus \{s\}} x_{(H,\pi,s),t} . \quad (120)$$

For any $s \in [r]$ and any graph $(H, \pi) \in \mathcal{H}_s$ as in part (b) of Definition 40 we obtain, using the definitions in line 10 and line 31,

$$w_{(H,\pi,s)}(v_1) \stackrel{(23)}{=} w_s(H, \pi) = w^{\hat{i}} = \sum_{t \in [r] \setminus \{s\}} d_t^{\hat{i}} \stackrel{(113b)}{=} \sum_{t \in [r] \setminus \{s\}} x_{(H,\pi,s),t} . \quad (121)$$

Together (120) and (121) prove the second part of the lemma. \square

For the next lemma, recall the definition of $\mathcal{C}(\mathcal{H}, F)$ in (22).

Lemma 43 (Graphs in $\mathcal{C}(\mathcal{H}_s, F)$ have smallest $d_\theta(\cdot)$ -value). *Let $\sigma \in [r]$ and $(H, \pi) \in \mathcal{H}_\sigma$, $\pi = (v_1, \dots, v_h)$, and let $s \in [r]$ and $(J, \tau) \in \mathcal{C}(\mathcal{H}_s, F)$, $\tau = (u_1, \dots, u_c)$. If $s = \sigma$, then we have*

$$d_\theta(J, u_1, w_{(J, \tau, \sigma)}) < x_{(H, \pi, \sigma), \sigma} = d_\theta(H, v_1, w_{(H, \pi, \sigma)}) .$$

If $s \neq \sigma$, then we have

$$d_\theta(J, u_1, w_{(J, \tau, s)}) \leq x_{(H, \pi, \sigma), s} .$$

Proof. Let i_{\max} denote the total number of iterations of the repeat-loop (*) in COMPUTEWEIGHTS().

First suppose that $s = \sigma$. Denoting by \hat{i} the largest index $\bar{i} \leq i_{\max}$ for which $\alpha_{\bar{i}} = \sigma$, the first part of Lemma 16 and the definitions in line 14 and line 23 show that $d_\theta(H, v_1, w_{(H, \pi, \sigma)}) \geq d_\sigma^{\hat{i}}$. By the termination condition in line 35 we also have $d_\theta(J, u_1, w_{(J, \tau, \sigma)}) < d_\sigma^{\hat{i}}$. We thus obtain $d_\theta(J, u_1, w_{(J, \tau, \sigma)}) < d_\theta(H, v_1, w_{(H, \pi, \sigma)})$. By the definition in (113) the right hand side of this last inequality equals $x_{(H, \pi, \sigma), \sigma}$, proving the first part of the lemma.

Now suppose that $s \neq \sigma$. By the definition in (113) we have

$$x_{(H, \pi, \sigma), s} = d_s^i \tag{122}$$

for some $1 \leq i \leq i_{\max}$. By the termination condition in line 36 and the observation that during the \bar{i} -th iteration of the repeat-loop (*), none of the families \mathcal{H}_t , $t \in [r] \setminus \{\alpha_{\bar{i}}\}$, is modified, we must have $\alpha_{i_{\max}} \neq s$, as we would have $\mathcal{H}_s = \mathcal{S}(F)$ otherwise, implying that $\mathcal{C}(\mathcal{H}_s, F)$ would be empty. So let \hat{i} be the largest index $\bar{i} \leq i_{\max} - 1$ for which $\alpha_{\bar{i}} = s$. By the definition in line 8 we have

$$d_\theta(J, u_1, w_{(J, \tau, s)}) \leq d_s^{\hat{i}+1} . \tag{123}$$

Using the first part of Lemma 16 twice we obtain

$$d_s^{\hat{i}+1} = \dots = d_s^{i_{\max}} \quad \text{and} \quad d_s^{i_{\max}} \leq d_s^i ,$$

which together with (122) and (123) yields the second part of the lemma. \square

Lemma 44 (Relation between $\lambda_\theta(\cdot)$ -value and x -point). *Let $s \in [r]$. For any graph $(H, \pi) \in \mathcal{H}_s$, $\pi = (v_1, \dots, v_h)$, we have*

$$\lambda_\theta(H, w_{(H, \pi, s)}) \geq 1 + \sum_{t \in [r]} x_{(H, \pi, s), t} .$$

Moreover, there is a subgraph $\hat{J} \subseteq H$ with $v_1 \in \hat{J}$ satisfying

$$\lambda_\theta(\hat{J}, w_{(\hat{J}, \pi|_{\hat{J}}, s)}) = 1 + \sum_{t \in [r]} x_{(\hat{J}, \pi|_{\hat{J}}, s), t} = 1 + \sum_{t \in [r]} x_{(H, \pi, s), t} .$$

Proof. We clearly have

$$\lambda_\theta(H, w_{(H, \pi, s)}) \stackrel{(19)}{=} \sum_{u \in H} (1 + w_{(H, \pi, s)}(u)) - e(H) \cdot \theta \stackrel{(18)}{\geq} 1 + w_{(H, \pi, s)}(v_1) + d_\theta(H, v_1, w_{(H, \pi, s)}) . \tag{124}$$

By Lemma 42 the right hand side of (124) equals $1 + \sum_{t \in [r]} x_{(H, \pi, s), t}$, proving the first part of the lemma.

Now consider a graph \hat{J} from the family

$$\arg \min_{J \subseteq H: v_1 \in J} \left(\sum_{u \in J \setminus v_1} (1 + w_{(H, \pi, s)}(u)) - e(J) \cdot \theta \right) . \tag{125}$$

Using the definition of $d_\theta()$ in (18) we obtain

$$d_\theta(H, v_1, w_{(H, \pi, s)}) \stackrel{(18), (125)}{=} \sum_{u \in \hat{J} \setminus v_1} (1 + w_{(H, \pi, s)}(u)) - e(\hat{J}) \cdot \theta \stackrel{(18), (125)}{=} d_\theta(\hat{J}, v_1, w_{(H, \pi, s)}) . \quad (126)$$

Furthermore, Lemma 24 yields that

$$w_{(H, \pi, s)}(u) = w_{(\hat{J}, \pi|_{\hat{J}}, s)}(u) \quad \text{for all } u \in \hat{J} . \quad (127)$$

Recall from the first part of the proof that the right hand side of (124) equals $1 + \sum_{t \in [r]} x_{(H, \pi, s), t}$. Applying (126), (127) and Lemma 42 shows that the right hand side of (124) also equals $1 + \sum_{t \in [r]} x_{(\hat{J}, \pi|_{\hat{J}}, s), t}$. Furthermore, applying the first equality in (126), (127) and the definition of $\lambda_\theta()$ in (19) shows that the right hand side of (124) equals $\lambda_\theta(\hat{J}, w_{(\hat{J}, \pi|_{\hat{J}}, s)})$, completing the proof of the second part of the lemma. \square

We say that a family \mathcal{D} of ordered graphs is *closed under taking subgraphs that contain the youngest vertex* if for any $(H, \pi) \in \mathcal{D}$, $\pi = (v_1, \dots, v_h)$, we have that for every $J \subseteq H$ with $v_1 \in J$ the ordered graph $(J, \pi|_J)$ is also contained in \mathcal{D} .

Note that the families \mathcal{D}_s , $s \in [r]$, used by the strategy PAINT() and defined in (106) are nonempty and closed under taking subgraphs that contain the youngest vertex.

Lemma 45 (*x-point of $\lambda_\theta()$ -minimizing graphs*). *Let $s \in [r]$ and $\mathcal{D}_s \subseteq \mathcal{H}_s$ a nonempty family of ordered graphs that is closed under taking subgraphs that contain the youngest vertex. For any graph (J, τ) from the family*

$$\arg \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)})$$

we have

$$\lambda_\theta(J, w_{(J, \tau, s)}) = 1 + \sum_{t \in [r]} x_{(J, \tau, s), t} .$$

Proof. The claim follows immediately from Lemma 44, using the closure property of the family \mathcal{D}_s and the choice of (J, τ) . \square

Lemma 46 (*$d_\theta()$ -value of $\lambda_\theta()$ -minimizing graphs*). *Let $s \in [r]$ and let $\mathcal{D}_s \subseteq \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$ be a nonempty family of ordered graphs that is closed under taking subgraphs that contain the youngest vertex. Furthermore, let (J, τ) , $\tau = (u_1, \dots, u_c)$, be an inclusion-minimal graph from the family*

$$\arg \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)}) .$$

Then we have

$$\lambda_\theta(J \setminus u_1, w_{(J, \tau, s)}) - \deg_J(u_1) \cdot \theta = d_\theta(J, u_1, w_{(J, \tau, s)}) . \quad (128)$$

Proof. We distinguish two cases depending on whether $(J, \tau) \in \mathcal{D}_s \subseteq \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$ is contained in \mathcal{H}_s or in $\mathcal{C}(\mathcal{H}_s, F)$.

If $(J, \tau) \in \mathcal{H}_s$, then by Lemma 45 we have

$$\lambda_\theta(J, w_{(J, \tau, s)}) = 1 + \sum_{t \in [r]} x_{(J, \tau, s), t} . \quad (129)$$

Rewriting the left hand side of (129) according to (20) and the right hand side according to Lemma 42 yields the desired equality (128).

We now consider the case $(J, \tau) \in \mathcal{C}(\mathcal{H}_s, F)$ (in this case we have $\lambda_\theta(J, w_{(J, \tau, s)}) = -\infty$ by Lemma 15). We clearly have

$$\lambda_\theta(J \setminus u_1, w_{(J, \tau, s)}) - \deg_J(u_1) \cdot \theta \stackrel{(19)}{=} \sum_{u \in J \setminus u_1} (1 + w_{(J, \tau, s)}(u)) - e(J) \cdot \theta \stackrel{(18)}{\geq} d_\theta(J, u_1, w_{(J, \tau, s)}) \quad , \quad (130)$$

and it remains to show that this inequality is in fact an equality. If the last inequality in (130) were strict, then, as in the proof of Lemma 44 (cf. (125), (126) and (127)), there would be a proper subgraph $\hat{J} \subsetneq J$ with $u_1 \in \hat{J}$ satisfying

$$d_\theta(J, u_1, w_{(J, \tau, s)}) = d_\theta(\hat{J}, u_1, w_{(\hat{J}, \tau|_{\hat{J}}, s)}) \quad . \quad (131)$$

As $(J, \tau) \in \mathcal{C}(\mathcal{H}_s, F)$ we have $(J \setminus u_1, \tau \setminus u_1) \in \mathcal{H}_s$, which by Lemma 22 implies that $(\hat{J} \setminus u_1, \tau|_{\hat{J} \setminus u_1}) \in \mathcal{H}_s$ as well. Hence $(\hat{J}, \tau|_{\hat{J}})$ must be in $\mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$. Using (131) and the first part of Lemma 43 shows that $(\hat{J}, \tau|_{\hat{J}})$ must be contained in $\mathcal{C}(\mathcal{H}_s, F)$. But then we have $\lambda_\theta(\hat{J}, w_{(\hat{J}, \tau|_{\hat{J}}, s)}) = -\infty$ by Lemma 15, a contradiction to the inclusion-minimality of (J, τ) (here we used again the closure property of the family \mathcal{D}_s). Therefore the last inequality in (130) holds with equality, proving the lemma also in this case. \square

Lemma 47 ($\lambda_\theta()$ -minimizing graphs in \mathcal{H}'_s). *Let $s \in [r]$ and let $\mathcal{D}_s \subseteq \mathcal{H}_s$ be a nonempty family of ordered graphs that is closed under taking subgraphs that contain the youngest vertex. Furthermore, let (J, τ) be an inclusion-minimal graph from the family*

$$\arg \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)})$$

and suppose that $x_{(J, \tau, s)}$ is the starting point of some s -segment of the walk \mathcal{W} defined in Lemma 41. Then (J, τ) is contained in \mathcal{H}'_s .

Proof. By Lemma 45 we have

$$\lambda_\theta(J, w_{(J, \tau, s)}) = 1 + \sum_{t \in [r]} x_{(J, \tau, s), t} \quad . \quad (132)$$

Let u_1 denote the youngest vertex of (J, τ) and let $\tilde{J} \subsetneq J$ be any proper subgraph of J with $u_1 \in \tilde{J}$. By Lemma 44 there is a subgraph $\hat{J} \subseteq \tilde{J}$ with $u_1 \in \hat{J}$ satisfying

$$\lambda_\theta(\hat{J}, w_{(\hat{J}, \tau|_{\hat{J}}, s)}) = 1 + \sum_{t \in [r]} x_{(\tilde{J}, \tau|_{\tilde{J}}, s), t} \quad . \quad (133)$$

By the inclusion-minimal choice of (J, τ) , (133) must be strictly larger than (132), i.e., we have

$$1 + \sum_{t \in [r]} x_{(J, \tau, s), t} < 1 + \sum_{t \in [r]} x_{(\tilde{J}, \tau|_{\tilde{J}}, s), t} \quad ,$$

in particular

$$x_{(J, \tau, s)} \neq x_{(\tilde{J}, \tau|_{\tilde{J}}, s)} \quad .$$

Using this observation together with the assumption that $x_{(J, \tau, s)}$ is the starting point of some s -segment of \mathcal{W} , it follows from property (iii) in Lemma 41 that (J, τ) must be contained in \mathcal{H}'_s . \square

Lemma 48 (x -points of graphs from \mathcal{H}'_s on the walk). *Let $s \in [r]$ and $(J, \tau) \in \mathcal{H}'_s$, $\tau = (u_1, \dots, u_c)$. Moreover, let $1 \leq b \leq c-1$ and define $(J^{-b}, \tau^{-b}) := (J \setminus \{u_1, \dots, u_b\}, \tau \setminus \{u_1, \dots, u_b\})$. Then $x_{(J, \tau, s)}$ is lower than $x_{(J^{-b}, \tau^{-b}, s)}$ on the walk \mathcal{W} defined in Lemma 41 and both points are contained in different s -segments of this walk.*

Proof. By the definition of \mathcal{H}'_s in (108) we have $(J, \tau) \in \mathcal{C}^{i,1}$ for some $i \geq 1$ with $\alpha_i = s$, i.e., (J, τ) was added to the family \mathcal{H}_s in the first iteration of the repeat-loop (**) in the i -th iteration of the repeat-loop (*) of COMPUTEWEIGHTS(). By the definition in (113a) we have

$$x_{(J, \tau, s), s} = d_\theta(J, u_1, w_{(J, \tau, s)}) = d_s^i . \quad (134)$$

By Lemma 22, the graph (J^{-b}, τ^{-b}) was added to the family \mathcal{H}_s either before the graph (J, τ) or together with it. But as (J^{-b}, τ^{-b}) is a predecessor of (J, τ) in the tree $\mathcal{T}(F)$ defined after (21), it follows from the definition in line 14 that (J^{-b}, τ^{-b}) must have already been contained in \mathcal{H}_s at the beginning of the i -th iteration of the repeat-loop (*). Applying Lemma 18 yields

$$x_{(J^{-b}, \tau^{-b}, s), s} \stackrel{(113)}{=} d_\theta(J^{-b}, u_{b+1}, w_{(J^{-b}, \tau^{-b}, s)}) > d_s^i . \quad (135)$$

Combining (134) and (135) shows that $x_{(J, \tau, s)}$ is lower than $x_{(J^{-b}, \tau^{-b}, s)}$ on the walk \mathcal{W} . As by the assumption $(J, \tau) \in \mathcal{H}'_s$ and property (iv) from Lemma 41 the point $x_{(J, \tau, s)}$ is the starting point of an s -segment of \mathcal{W} , this implies that both points must be contained in different s -segments of this walk. \square

5.4. Analysis of PAINT(). We are now in a position to actually analyze our Painter strategy PAINT(). Recall from Section 5.1 that the parameter $d(s)$ defined in (107) might be equal to $-\infty$ for some colors $s \in [r]$ (intuitively, Painter considers such a color extremely dangerous). The following lemma shows that PAINT() never chooses such a color.

Lemma 49 (Painter strategy creates only graphs from \mathcal{H}_s). *Consider a fixed step of the game, and let the families $\mathcal{D}_s \subseteq \mathcal{S}(F)$, $s \in [r]$, and the values $d(s) \in \mathbb{R} \cup \{-\infty\}$ be defined as in (106) and (107), respectively. For any $\sigma \in \arg \max_{s \in [r]} d(s)$ the value $d(\sigma)$ is finite and we have $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$.*

Consequently, playing according to the strategy PAINT() throughout ensures that for all $s \in [r]$ we always have $\mathcal{D}_s \subseteq \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$ (even if ties are broken arbitrarily).

Proof. By the definition in (107) and Lemma 15 the value $d(s)$ is finite if and only if $\mathcal{D}_s \subseteq \mathcal{H}_s$. By the termination condition in line 36 there is some color $s \in [r]$ for which $\mathcal{H}_s = \mathcal{S}(F)$. For this color we therefore have $\mathcal{D}_s \subseteq \mathcal{H}_s$, implying that the corresponding value $d(s)$ is finite. This shows that for any $\sigma \in \arg \max_{s \in [r]} d(s)$, the value $d(\sigma)$ is finite and therefore $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$, proving the first part of the lemma. The second part follows inductively by observing that the strategy PAINT() in each step picks a color $\sigma \in \arg \max_{s \in [r]} d(s)$ (regardless of the tie-breaking rule), showing that in this step only graphs from the family $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$ in color σ are created on the board. \square

The following lemma shows that the tie-breaking rule of the strategy PAINT(), which uses the families \mathcal{H}'_s and \mathcal{J}_s defined in (108) and (109), is indeed well-defined.

Lemma 50 (Well-definedness of Painter strategy). *Ties in the strategy PAINT() can arise only between two different colors, and if they arise then for exactly one of the two colors (say σ) we have $\mathcal{J}_\sigma \cap \mathcal{H}'_\sigma = \emptyset$, and for the other color (say s) we have $\mathcal{J}_s \cap \mathcal{H}'_s \neq \emptyset$. (Thus the tie-breaking rule will decide for color σ .)*

If such a tie arises, the walk \mathcal{W} defined in Lemma 41 contains a σ -segment Γ whose endpoint $x \in \mathbb{R}^r$ is also the starting point of an s -segment Γ' , and for any $(J_\sigma, \tau_\sigma) \in \mathcal{J}_\sigma$ and any $(J_s, \tau_s) \in \mathcal{J}_s$ we have $x_{(J_\sigma, \tau_\sigma, \sigma)} = x_{(J_s, \tau_s, s)} = x$.

Proof. Recall that the families \mathcal{D}_s , $s \in [r]$, defined in (106) are nonempty and closed under taking subgraphs that contain the youngest vertex. Fix some color $\sigma \in [r]$ such that

$$d(s) \leq d(\sigma) \quad , \quad s \in [r] \setminus \{\sigma\} , \quad (136)$$

for the values $d(s)$, $s \in [r]$, defined in (107). The tie-breaking rule of the strategy `PAINT()` is only considered if the inequality in (136) is tight for some color different from σ . We fix such a color $s \in [r] \setminus \{\sigma\}$ for which

$$d(s) = d(\sigma) . \quad (137)$$

The first part of Lemma 49 yields with (136) and (137) that $\mathcal{D}_s \subseteq \mathcal{H}_s$ and $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$ (and that $d(s)$ and $d(\sigma)$ are finite values). Thus by the definition in (109) we also have $\mathcal{J}_s \subseteq \mathcal{H}_s$ and $\mathcal{J}_\sigma \subseteq \mathcal{H}_\sigma$. Fix some $(J_s, \tau_s) \in \mathcal{J}_s$ and some $(J_\sigma, \tau_\sigma) \in \mathcal{J}_\sigma$. By the definition in (109) and Lemma 45 we have

$$\lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) = 1 + \sum_{t \in [r]} x_{(J_s, \tau_s, s), t} \quad \text{and} \quad \lambda_\theta(J_\sigma, w_{(J_\sigma, \tau_\sigma, \sigma)}) = 1 + \sum_{t \in [r]} x_{(J_\sigma, \tau_\sigma, \sigma), t} . \quad (138)$$

Furthermore, using (137) and the definitions in (107) and (109) shows that

$$\lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) = \lambda_\theta(J_\sigma, w_{(J_\sigma, \tau_\sigma, \sigma)}) . \quad (139)$$

Combining (138) and (139) we obtain

$$1 + \sum_{t \in [r]} x_{(J_s, \tau_s, s), t} = 1 + \sum_{t \in [r]} x_{(J_\sigma, \tau_\sigma, \sigma), t} . \quad (140)$$

Note that the points $x_{(J_s, \tau_s, s)}$ and $x_{(J_\sigma, \tau_\sigma, \sigma)}$ are elements of the walk \mathcal{W} defined in Lemma 41. By Lemma 38 the relation (140) implies that $x_{(J_s, \tau_s, s)} = x_{(J_\sigma, \tau_\sigma, \sigma)}$, i.e., the graphs (J_s, τ_s) and (J_σ, τ_σ) (and all other graphs in the families \mathcal{J}_s and \mathcal{J}_σ) have the same associated point on the walk \mathcal{W} . By property (i) in Lemma 41, $x_{(J_s, \tau_s, s)}$ is contained in an s -segment and $x_{(J_\sigma, \tau_\sigma, \sigma)}$ in a σ -segment of \mathcal{W} , implying that $x_{(J_s, \tau_s, s)} = x_{(J_\sigma, \tau_\sigma, \sigma)}$ must be the endpoint of some segment and the starting point of the next lower segment. As on the walk \mathcal{W} , only pairs of consecutive segments have a point in common, this shows that the inequality (136) can be tight for at most one color different from σ , proving that ties can arise only between two different colors.

Assume w.l.o.g. that for all $(J_\sigma, \tau_\sigma) \in \mathcal{J}_\sigma$, the point $x_{(J_\sigma, \tau_\sigma, \sigma)}$ is the endpoint of some σ -segment Γ and for all $(J_s, \tau_s) \in \mathcal{J}_s$ the point $x_{(J_s, \tau_s, s)}$ is the starting point of the next lower s -segment Γ' . By property (iv) from Lemma 41 we have $\mathcal{J}_\sigma \cap \mathcal{H}'_\sigma = \emptyset$, and by Lemma 47 we have $\mathcal{J}_s \cap \mathcal{H}'_s \neq \emptyset$. This proves the first part of the lemma and shows that our tie-breaking rule is well-defined.

Note that the segment Γ is higher on the walk \mathcal{W} than Γ' , and that the tie-breaking rule decides for the color σ corresponding to the higher of the two segments. Together with our previous observations about the location of the points $x_{(J_\sigma, \tau_\sigma, \sigma)}$ for all $(J_\sigma, \tau_\sigma) \in \mathcal{J}_\sigma$ and $x_{(J_s, \tau_s, s)}$ for all $(J_s, \tau_s) \in \mathcal{J}_s$ this proves the second part of the lemma. \square

The following lemma will be the key to proving Lemma 35, our main strategy invariant based on witness graphs.

Lemma 51 (Painter strategy ensures sufficient weight). *There is a constant $\varepsilon = \varepsilon(F, r, \theta, \alpha) > 0$ such that the following holds: Let $\sigma \in [r]$ denote the color selected by the strategy `PAINT()` in a certain step of the game given the families \mathcal{D}_s , $s \in [r]$, defined in (106). For every $s \in [r] \setminus \{\sigma\}$, let (J_s, τ_s) be an inclusion-minimal graph from the family*

$$\mathcal{J}_s \stackrel{(109)}{=} \arg \min_{(H, \pi) \in \mathcal{D}_s} \lambda_\theta(H, w_{(H, \pi, s)}) ,$$

and let u_{s1} denote the youngest vertex of (J_s, τ_s) .

Then for any graph $(H, \pi) \in \mathcal{D}_\sigma$, $\pi = (v_1, \dots, v_h)$, we have

$$\sum_{s \in [r] \setminus \{\sigma\}} (\lambda_\theta(J_s \setminus u_{s1}, w_{(J_s, \tau_s, s)}) - \deg_{J_s}(u_{s1})) \leq w_{(H, \pi, \sigma)}(v_1) . \quad (141)$$

If the inequality (141) is strict, then the difference between the right and left hand side is at least ε .

If on the other hand the inequality (141) is tight, then for every $s \in [r] \setminus \{\sigma\}$ we have $(J_s, \tau_s) \in \mathcal{H}'_s \cup \mathcal{C}(\mathcal{H}_s, F)$. Moreover, denoting by \mathcal{W} the walk defined in Lemma 41 and by Γ the σ -segment containing $x_{(H, \pi, \sigma)}$ on this walk, we have the following: if $(J_s, \tau_s) \in \mathcal{H}'_s$, then $x_{(J_s, \tau_s, s)}$ is the starting point of the next s -segment on \mathcal{W} that is lower than Γ , whereas if $(J_s, \tau_s) \in \mathcal{C}(\mathcal{H}_s, F)$, then there is no s -segment on \mathcal{W} lower than Γ .

Proof. Let

$$\varepsilon = \varepsilon(F, r, \theta, \alpha) := \min\{ |d_\theta(H, v_1, w_{(H, \pi, s)}) - d_\theta(J, u_1, w_{(J, \tau, s)})| \mid s \in [r] \wedge (H, \pi = (v_1, \dots, v_h)), (J, \tau = (u_1, \dots, u_c)) \in \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F) \wedge d_\theta(H, v_1, w_{(H, \pi, s)}) \neq d_\theta(J, u_1, w_{(J, \tau, s)}) \} > 0 \quad (142)$$

(recall that for all $s \in [r]$ the $d_\theta()$ -value of all graphs in $\mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$ is a finite real number). Note that in the boundary case that for all $s \in [r]$ and all $(H, \pi) \in \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$, $\pi = (v_1, \dots, v_h)$, the value $d_\theta(H, v_1, w_{(H, \pi, s)})$ is the same (i.e., the walk \mathcal{W} defined in Lemma 41 degenerates to a single point), the minimum in (142) is over an empty set. We will see that in this case the inequality (141) is never strict. Therefore we may set ε to an arbitrary positive constant in this case, $\varepsilon := 1$, say.

Recall that the families \mathcal{D}_s , $s \in [r]$, are nonempty and closed under taking subgraphs that contain the youngest vertex. By the second part of Lemma 49 we have $\mathcal{D}_s \subseteq \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$ for all $s \in [r]$.

We first prove that (141) holds. By the definition of the strategy, the selected color $\sigma \in [r]$ satisfies

$$d(s) \leq d(\sigma) \quad , \quad s \in [r] \setminus \{\sigma\} \quad (143)$$

for the values $d(s)$, $s \in [r]$, defined in (107). Let (J_σ, τ_σ) be an arbitrary graph from the family

$$\mathcal{J}_\sigma \stackrel{(109)}{=} \arg \min_{(H, \pi) \in \mathcal{D}_\sigma} \lambda_\theta(H, w_{(H, \pi, \sigma)}) \quad .$$

By the definition in (107) and the choice of (J_s, τ_s) , $s \in [r]$, we have

$$\lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) = d(s) \quad , \quad s \in [r] \quad . \quad (144)$$

Combining (143) and (144) yields

$$\lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) \leq \lambda_\theta(J_\sigma, w_{(J_\sigma, \tau_\sigma, \sigma)}) \quad , \quad s \in [r] \setminus \{\sigma\} \quad . \quad (145)$$

By (143) and the first part of Lemma 49 we have $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$. We fix some graph $(H, \pi) \in \mathcal{D}_\sigma$, $\pi = (v_1, \dots, v_h)$. By Lemma 44 there is a subgraph $\hat{J} \subseteq H$ with $v_1 \in \hat{J}$ satisfying

$$\lambda_\theta(\hat{J}, w_{(\hat{J}, \pi|_{\hat{J}}, \sigma)}) = 1 + \sum_{t \in [r]} x_{(H, \pi, \sigma), t} \quad . \quad (146)$$

By the closure property of the family \mathcal{D}_σ and the choice of (J_σ, τ_σ) we have

$$\lambda_\theta(J_\sigma, w_{(J_\sigma, \tau_\sigma, \sigma)}) \leq \lambda_\theta(\hat{J}, w_{(\hat{J}, \pi|_{\hat{J}}, \sigma)}) \quad . \quad (147)$$

By Lemma 45 we have for every $s \in [r] \setminus \{\sigma\}$ for which (J_s, τ_s) is contained in \mathcal{H}_s that

$$\lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) = 1 + \sum_{t \in [r]} x_{(J_s, \tau_s, s), t} \quad . \quad (148)$$

For those $s \in [r] \setminus \{\sigma\}$ we thus obtain

$$1 + \sum_{t \in [r]} x_{(J_s, \tau_s, s), t} \stackrel{(148)}{=} \lambda_\theta(J_s, w_{(J_s, \tau_s, s)}) \stackrel{(145)}{\leq} \lambda_\theta(J_\sigma, w_{(J_\sigma, \tau_\sigma, \sigma)}) \stackrel{(146), (147)}{\leq} 1 + \sum_{t \in [r]} x_{(H, \pi, \sigma), t} \quad . \quad (149)$$

Note that if $(J_s, \tau_s) \in \mathcal{H}_s$, then $x_{(J_s, \tau_s, s)}$ is an element of the walk \mathcal{W} defined in Lemma 41 (the point $x_{(H, \pi, \sigma)}$ is clearly also an element of this walk as $\mathcal{D}_\sigma \subseteq \mathcal{H}_\sigma$). Using (149) and Lemma 38 yields that for every $s \in [r] \setminus \{\sigma\}$ for which (J_s, τ_s) is contained in \mathcal{H}_s we have

$$x_{(J_s, \tau_s, s), t} \leq x_{(H, \pi, \sigma), t} \quad \text{for all } t \in [r] , \quad (150)$$

from which we conclude using

$$x_{(J_s, \tau_s, s), s} \stackrel{(113)}{=} d_\theta(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) \quad (151)$$

that

$$d_\theta(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) \leq x_{(H, \pi, \sigma), s} . \quad (152)$$

For all $s \in [r] \setminus \{\sigma\}$ for which (J_s, τ_s) is not contained in \mathcal{H}_s but in $\mathcal{C}(\mathcal{H}_s, F)$, the relation (152) follows from the second part of Lemma 43.

Combining our previous observations and applying Lemma 42 and Lemma 46, we thus obtain

$$\begin{aligned} \sum_{s \in [r] \setminus \{\sigma\}} (\lambda_\theta(J_s \setminus u_{s1}, w_{(J_s, \tau_s, s)}) - \deg_{J_s}(u_{s1})) &\stackrel{(128)}{=} \sum_{s \in [r] \setminus \{\sigma\}} d_\theta(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) \\ &\stackrel{(152)}{\leq} \sum_{s \in [r] \setminus \{\sigma\}} x_{(H, \pi, \sigma), s} \stackrel{(119)}{=} w_{(H, \pi, \sigma)}(v_1) , \end{aligned} \quad (153)$$

proving (141).

If the inequality (153) is strict, then by (152) we have

$$d_\theta(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) < x_{(H, \pi, \sigma), s} \quad \text{for some } s \in [r] \setminus \{\sigma\} . \quad (154)$$

By the definition in (113) and the definition in line 8, the right hand side of (154) equals $d_\theta(\bar{J}, \bar{u}_1, w_{(\bar{J}, \bar{\tau}, s)})$ for some $(\bar{J}, \bar{\tau}) \in \mathcal{H}_s \cup \mathcal{C}(\mathcal{H}_s, F)$, where \bar{u}_1 denotes the youngest vertex of $(\bar{J}, \bar{\tau})$. With the definition in (142) it follows that the difference between the right and left hand side of (154) and therefore also the difference between the right and left hand side of (153) is at least ε .

If the inequality in (153) is tight, then by (152) we have

$$d_\theta(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) = x_{(H, \pi, \sigma), s} \quad \text{for all } s \in [r] \setminus \{\sigma\} . \quad (155)$$

Let Γ denote the σ -segment on the walk \mathcal{W} that contains the point $x_{(H, \pi, \sigma)}$. We fix some $s \in [r] \setminus \{\sigma\}$ and distinguish the cases whether (J_s, τ_s) is contained in \mathcal{H}_s or in $\mathcal{C}(\mathcal{H}_s, F)$.

We first consider the case that $(J_s, \tau_s) \in \mathcal{H}_s$. We claim that the s -segment Γ' on the walk \mathcal{W} that contains the point $x_{(J_s, \tau_s, s)}$ is lower on the walk \mathcal{W} than Γ : This is trivially true if one of the inequalities in (150) is strict. If on the other hand (150) holds with equality for all $t \in [r]$, then also all inequalities in (149) are tight, from which we conclude using (144) that $d(s) = d(\sigma)$, i.e., we have a tie between the colors s and σ . In this case, by the second part of Lemma 50, our tie-breaking rule ensures that Γ' is lower on the walk \mathcal{W} than Γ . From (151) and (155) it follows that Γ' must be the next s -segment on \mathcal{W} that is lower than Γ and that $x_{(J_s, \tau_s, s)}$ must be the starting point of Γ' . Applying Lemma 47 shows that $(J_s, \tau_s) \in \mathcal{H}'_s$ (note the inclusion-minimal choice of (J_s, τ_s)), completing the proof in this case.

It remains to consider the case $(J_s, \tau_s) \in \mathcal{C}(\mathcal{H}_s, F)$. Suppose for the sake of contradiction that there was some s -segment Γ' that is lower than Γ on the walk \mathcal{W} . By property (v) in Lemma 41 the segment Γ' can not be lowest segment of \mathcal{W} , as we would otherwise have $\mathcal{H}_s = \mathcal{S}(F)$ and therefore $\mathcal{C}(\mathcal{H}_s, F) = \emptyset$. So Γ' has an endpoint $x \in \mathbb{R}^r$ which clearly satisfies

$$x_s < x_{(H, \pi, \sigma), s} \quad (156)$$

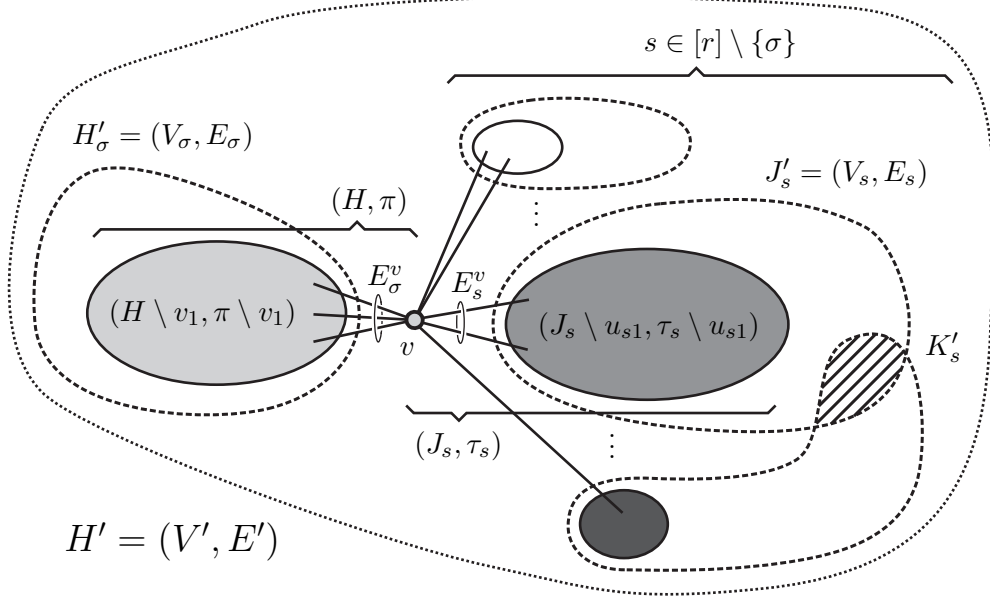


FIGURE 6. Notations used in the first part of the proof of Lemma 35.

and which is also the starting point of the next lower segment $\bar{\Gamma}$ (the segment $\bar{\Gamma}$ is an \bar{s} -segment for some $\bar{s} \in [r] \setminus \{s\}$). By property (ii) of Lemma 41 there is some $(\bar{J}, \bar{\tau}) \in \mathcal{H}_{\bar{s}}$ such that

$$x_{(\bar{J}, \bar{\tau}, \bar{s})} = x . \quad (157)$$

Applying the second part of Lemma 43 we thus obtain

$$d_{\theta}(J_s, u_{s1}, w_{(J_s, \tau_s, s)}) \leq x_{(\bar{J}, \bar{\tau}, \bar{s}), s} \stackrel{(157)}{=} x_s \stackrel{(156)}{<} x_{(H, \pi, \sigma), s} ,$$

contradicting (155). This completes the proof also in this case. \square

5.5. Proof of Lemma 35. We are now ready to prove Lemma 35, our main strategy invariant.

Proof of Lemma 35. For the reader's convenience, Figure 6 illustrates the notations used in the first part of the proof.

Let

$$v_{\max} = v_{\max}(F, r, \theta, \alpha) := r^{v(F)/\varepsilon + (v(F)/\varepsilon + 1)(r \cdot |S(F)| + 1)(v(F) + 1) + 2} \cdot v(F) + 1 , \quad (158)$$

where $\varepsilon = \varepsilon(F, r, \theta, \alpha)$ is the constant guaranteed by Lemma 51 (and explicitly defined in (142)).

We argue by induction over the number of vertices of the board. For the induction base consider the board at the beginning of the game when no vertex is added yet. It is convenient for the proof to extend the statement of the lemma to (H, π) being the null graph (the graph whose vertex set is empty). For this graph we define H' to be the null graph as well. Clearly, for every $s \in [r]$, every copy of the null graph (H, π) ‘in color s ’ on the board is contained in this subgraph H' of the board, and we have $\mu_{\theta}(H') = 0 = \lambda_{\theta}(H, w_{(H, \pi, s)})$ and $v(H') = 0 \leq v_{\max}$. This shows that the second condition of the lemma holds at the beginning of the game and settles the induction base.

For the induction step, let v denote the vertex added in the current step of the game, \mathcal{D}_s , $s \in [r]$, the families defined in (106), and σ the color the strategy $\text{PAINT}()$ assigns to the vertex v . By the first part of Lemma 49, we have $\mathcal{D}_{\sigma} \subseteq \mathcal{H}_{\sigma}$.

For a fixed graph $(H, \pi) \in \mathcal{D}_\sigma$, $\pi = (v_1, \dots, v_h)$, we consider a fixed copy of $(H \setminus v_1, \pi \setminus v_1)$ in color σ that is completed by v to a copy of (H, π) in this color. Denoting by E_σ^v the corresponding set of edges incident to v , we clearly have

$$|E_\sigma^v| = \deg_H(v_1) . \quad (159)$$

By induction, we know that this copy of $(H \setminus v_1, \pi \setminus v_1)$ is contained in a graph $H'_\sigma = (V_\sigma, E_\sigma)$ with

$$\mu_\theta(H'_\sigma) \leq \lambda_\theta(H \setminus v_1, w_{(H, \pi, \sigma)}) \quad (160)$$

(recall from (23) that $w_{(H, \pi, \sigma)}(u) = w_{(H \setminus v_1, \pi \setminus v_1, \sigma)}(u)$ for all $u \in H \setminus v_1$) and

$$v(H'_\sigma) \leq v_{\max} . \quad (161)$$

For every $s \in [r] \setminus \{\sigma\}$, let (J_s, τ_s) be an inclusion-minimal graph from the family $\mathcal{J}_s \subseteq \mathcal{D}_s$ defined in (109), and let u_{s1} denotes the youngest vertex of (J_s, τ_s) . For each $s \in [r] \setminus \{\sigma\}$ we consider a fixed copy of $(J_s \setminus u_{s1}, \tau_s \setminus u_{s1})$ in color s that is completed by v to a copy of (J_s, τ_s) (the vertex v has the color $\sigma \neq s$, so the resulting copy is not monochromatic). Denoting by E_s^v the corresponding set of edges incident to v , we clearly have

$$|E_s^v| = \deg_{J_s}(u_{s1}) \quad , \quad s \in [r] \setminus \{\sigma\} . \quad (162)$$

By induction, those copies of $(J_s \setminus u_{s1}, \tau_s \setminus u_{s1})$ are contained in graphs $J'_s = (V_s, E_s)$ with

$$\mu_\theta(J'_s) \leq \lambda_\theta(J_s \setminus u_{s1}, w_{(J_s, \tau_s, s)}) \quad , \quad s \in [r] \setminus \{\sigma\} , \quad (163)$$

and

$$v(J'_s) \leq v_{\max} \quad , \quad s \in [r] \setminus \{\sigma\} . \quad (164)$$

Applying Lemma 51 shows that the graphs (H, π) and (J_s, τ_s) , $s \in [r] \setminus \{\sigma\}$, satisfy

$$\sum_{s \in [r] \setminus \{\sigma\}} (\lambda_\theta(J_s \setminus u_{s1}, w_{(J_s, \tau_s, s)}) - \deg_{J_s}(u_{s1})) \leq w_{(H, \pi, \sigma)}(v_1) . \quad (165)$$

If $\mu_\theta(H'_\sigma) < 0$ or $\mu_\theta(J'_s) < 0$ for some $s \in [r] \setminus \{\sigma\}$, we have found a graph K' with $\mu_\theta(K') < 0$ and $v(K') \leq v_{\max}$ (see (161) and (164)). Otherwise we have $\mu_\theta(H'_\sigma) \geq 0$ and $\mu_\theta(J'_s) \geq 0$ for all $s \in [r] \setminus \{\sigma\}$. We will argue later that this implies even stronger bounds on the number of vertices of H'_σ and J'_s , namely

$$\begin{aligned} v(H'_\sigma) &\leq (v_{\max} - 1)/r , \\ v(J'_s) &\leq (v_{\max} - 1)/r \quad , \quad s \in [r] \setminus \{\sigma\} . \end{aligned} \quad (166)$$

We define the graph $H' = (V', E')$ as

$$\begin{aligned} V' &:= \{v\} \cup \bigcup_{s \in [r]} V_s , \\ E' &:= \bigcup_{s \in [r]} (E_s \cup E_s^v) \end{aligned} \quad (167)$$

(see Figure 6). This graph clearly contains the copy of (H, π) in color σ we are considering.

Furthermore, we define for $2 \leq s \leq r$ the graphs

$$K'_s := \left(V_s \cap \bigcup_{1 \leq t \leq s-1} V_t, E_s \cap \bigcup_{1 \leq t \leq s-1} E_t \right) . \quad (168)$$

From (166) and (168) we conclude that $v(K'_s) \leq (v_{\max} - 1)/r \leq v_{\max}$, $2 \leq s \leq r$. Therefore, if $\mu_\theta(K'_s) < 0$ for some $2 \leq s \leq r$, then we have found a graph K' with $\mu_\theta(K') < 0$ and $v(K') \leq v_{\max}$. Otherwise we have

$$\mu_\theta(K'_s) \geq 0 \quad , \quad 2 \leq s \leq r . \quad (169)$$

With (159) and (162) we obtain from (167) and (168) that

$$\begin{aligned} v(H') &= 1 + v(H'_\sigma) + \sum_{s \in [r] \setminus \{\sigma\}} v(J'_s) - \sum_{2 \leq s \leq r} v(K'_s) , \\ e(H') &= e(H'_\sigma) + \deg_H(v_1) + \sum_{s \in [r] \setminus \{\sigma\}} (e(J'_s) + \deg_{J_s}(u_{s1})) - \sum_{2 \leq s \leq r} e(K'_s) . \end{aligned} \quad (170)$$

Combining our previous observations yields

$$\begin{aligned} \mu_\theta(H') &\stackrel{(6),(170)}{=} 1 + \mu_\theta(H'_\sigma) - \deg_H(v_1) \cdot \theta + \sum_{s \in [r] \setminus \{\sigma\}} (\mu_\theta(J'_s) - \deg_{J_s}(u_{s1}) \cdot \theta) - \sum_{2 \leq s \leq r} \mu_\theta(K'_s) \\ &\stackrel{(160),(163),(169)}{\leq} 1 + \lambda_\theta(H \setminus v_1, w_{(H,\pi,\sigma)}) - \deg_H(v_1) \cdot \theta + \underbrace{\sum_{s \in [r] \setminus \{\sigma\}} (\lambda_\theta(J_s \setminus u_{s1}, w_{(J_s,\tau_s,s)}) - \deg_{J_s}(u_{s1}) \cdot \theta)}_{\stackrel{(165)}{\leq} w_{(H,\pi,\sigma)}(v_1)} \\ &\leq \lambda_\theta(H \setminus v_1, w_{(H,\pi,\sigma)}) + 1 + w_{(H,\pi,\sigma)}(v_1) - \deg_H(v_1) \cdot \theta \stackrel{(20)}{=} \lambda_\theta(H, w_{(H,\pi,\sigma)}) \end{aligned} \quad (171)$$

which proves (111). From (166) and (167) we conclude that $v(H') \leq v_{\max}$.

It remains to show (166), i.e. that for every graph H' as defined in (167) with $\mu_\theta(H') \geq 0$ we have $v(H') \leq (v_{\max} - 1)/r$. For the reader's convenience, the notations used in this part of the proof are illustrated in Figure 7.

In the above argument we constructed the graph H' containing the copy of (H, π) in color σ inductively from the graph H'_σ containing the copy of $(H \setminus v_1, \pi \setminus v_1)$ in color σ and the graphs J'_s containing the copies of $(J_s \setminus u_{s1}, \tau_s \setminus u_{s1})$ in color s , $s \in [r] \setminus \{\sigma\}$. We associate this inductive construction with a node-colored rooted tree $\mathcal{T}(H')$, some of whose non-leaf nodes receive a special marking (we refer to it as a *flag*), as follows (see the upper part of Figure 7): The nodes of $\mathcal{T}(H')$ correspond to monochromatic copies of graphs from $\mathcal{S}(F)$ on the board (the same copy may appear as a node multiple times). If (H, π) is the null graph ‘in color σ ’ (recall that in this case H' is the null graph as well), $\mathcal{T}(H')$ consists only of this copy of (H, π) as an isolated node which receives the color σ . Otherwise $\mathcal{T}(H')$ consists of the copy of (H, π) as the root node joined to r subtrees, $\mathcal{T}(H'_\sigma)$ and $\mathcal{T}(J'_s)$ for all $s \in [r] \setminus \{\sigma\}$. The root node receives the color σ , and it is flagged if and only if the instance of the inequality (165) corresponding to this induction step is strict. Note that the tree $\mathcal{T}(H')$ captures only the logical structure of the inductive history of H' . Overlappings (captured by the graphs K'_s , $2 \leq s \leq r$) are completely neglected.

Every flagged node of $\mathcal{T}(H')$ corresponds to a strict inequality in (165). In this case, inequality (171) is strict as well, with a difference of at least ε between the right and left hand side, where ε is the constant guaranteed by Lemma 51. Consequently, each flagged node contributes a term of $-\varepsilon$ to the right hand side of (171) in the corresponding induction step. Accumulating these terms along the induction yields that

$$\mu_\theta(H') \leq \lambda_\theta(H, w_{(H,\pi,\sigma)}) - f(H') \cdot \varepsilon , \quad (172)$$

where $f(H')$ denotes the number of flagged nodes in $\mathcal{T}(H')$.

By Lemma 15 we have $\lambda_\theta(H, w_{(H,\pi,\sigma)}) \leq v(F)$. Thus if $\mu_\theta(H') \geq 0$, then by (172) the tree $\mathcal{T}(H')$ has at most $\lambda_\theta(H, w_{(H,\pi,\sigma)})/\varepsilon \leq v(F)/\varepsilon$ many flagged nodes. We will show that this bound on the number of flagged nodes of $\mathcal{T}(H')$ implies the claimed bound of $(v_{\max} - 1)/r$ on the number of vertices of H' . To that end, we first show that the length of any descending path in $\mathcal{T}(H')$ that consists only of non-flagged nodes is bounded by a constant depending only on F and r . We will

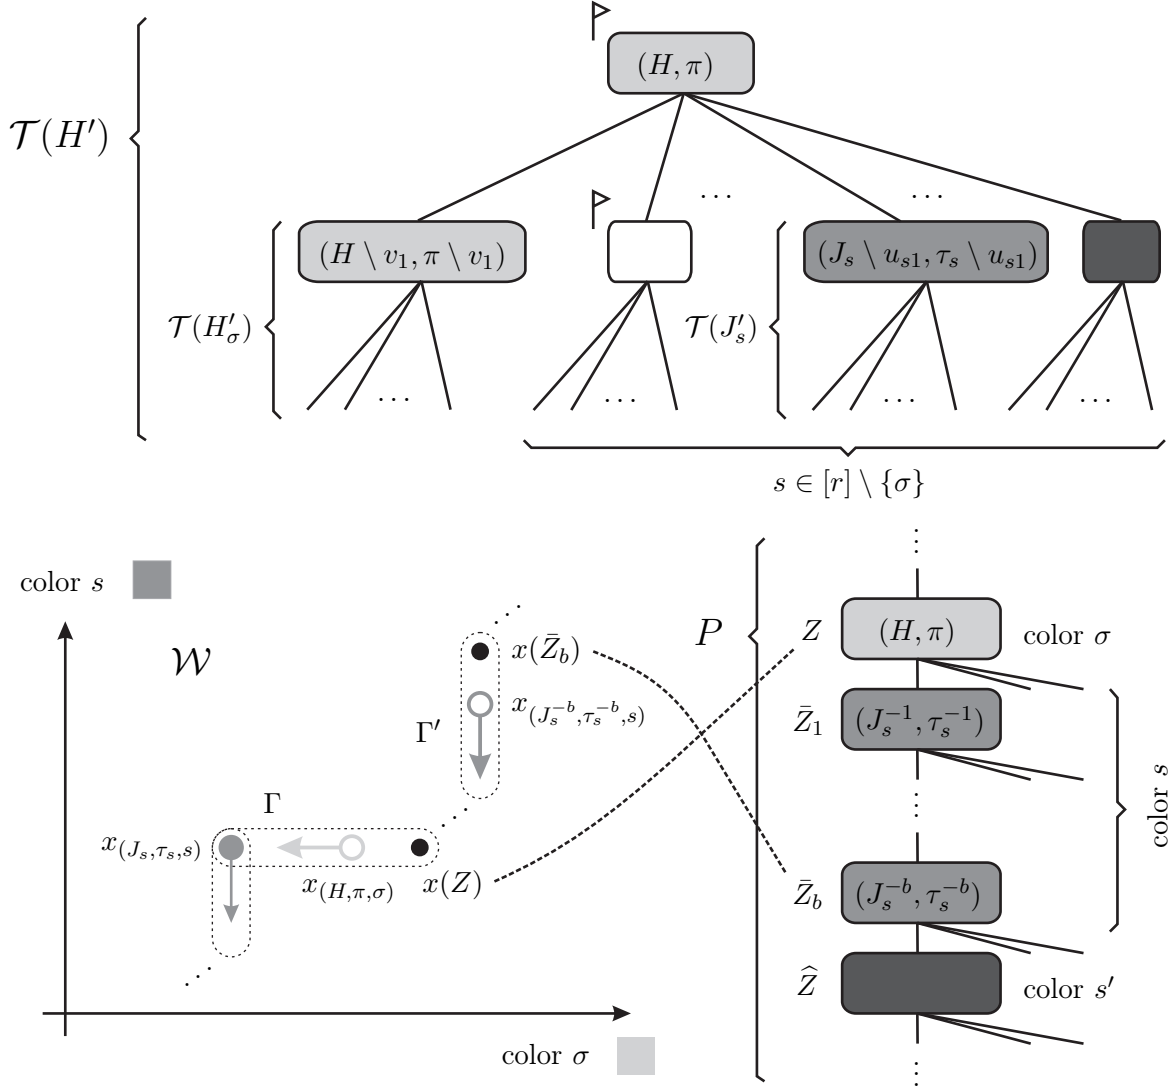


FIGURE 7. Notations used in the second part of the proof of Lemma 35.

do so by showing that any descending sequence of non-flagged nodes in $\mathcal{T}(H')$ corresponds to an ascending sequence of points on the walk \mathcal{W} defined in Lemma 41.

Specifically, we assign to every non-leaf node Z in $\mathcal{T}(H')$ a point $x(Z) \in \mathbb{R}^r$ on the walk \mathcal{W} as follows: Let σ denote the color of Z and (H, π) the graph for which a copy in color σ is represented by the node Z . We define $x(Z)$ as the starting point of the σ -segment that contains the point $x_{(H, \pi, \sigma)}$ on the walk \mathcal{W} .

Consider a descending sequence $Z, \bar{Z}_1, \dots, \bar{Z}_b, \hat{Z}$ of consecutive non-flagged nodes in $\mathcal{T}(H')$, where Z has some color $\sigma \in [r]$, all nodes $\bar{Z}_1, \dots, \bar{Z}_b$ have the same color $s \in [r] \setminus \{\sigma\}$ and \hat{Z} has some color $s' \in [r] \setminus \{s\}$ (see the lower right part of Figure 7). Let (H, π) be the graph for which a copy in color σ is represented by the node Z , and (J_s, τ_s) , $\tau_s = (u_{s1}, \dots, u_{sc})$, the graph for which a copy of $(J_s \setminus u_{s1}, \tau_s \setminus u_{s1})$ in color s is represented by the node \bar{Z}_1 . Using these definitions, clearly the nodes $\bar{Z}_1, \dots, \bar{Z}_b$ represent a sequence of nested copies of (J_s^{-a}, τ_s^{-a}) , $a = 1, \dots, b$, in color s , where $(J_s^{-a}, \tau_s^{-a}) := (J_s \setminus \{u_{s1}, \dots, u_{sa}\}, \tau_s \setminus \{u_{s1}, \dots, u_{sa}\})$. Let Γ denote the σ -segment of the walk \mathcal{W} that contains the point $x_{(H, \pi, \sigma)}$ and Γ' the s -segment that contains the point $x_{(J_s^{-b}, \tau_s^{-b}, s)}$

(the starting points of these segments are $x(Z)$ and $x(\bar{Z}_b)$, respectively; see the lower left part of Figure 7). As the node Z is not flagged, the corresponding instance of the inequality (165) is tight. Hence, by Lemma 51, we have $(J_s, \tau_s) \in \mathcal{H}'_s \cup \mathcal{C}(\mathcal{H}_s, F)$, and if $(J_s, \tau_s) \in \mathcal{H}'_s$, then $x_{(J_s, \tau_s, s)}$ is the starting point of the next s -segment on \mathcal{W} that is lower than Γ , whereas if $(J_s, \tau_s) \in \mathcal{C}(\mathcal{H}_s, F)$, then there is no s -segment on \mathcal{W} lower than Γ . In the first case we apply Lemma 48 to conclude that Γ is lower than Γ' on the walk \mathcal{W} , in the second case this conclusion is trivially true. It follows that in any case $x(Z)$ is lower on the walk than $x(\bar{Z}_b)$.

Let now P be a descending path in $\mathcal{T}(H')$ that consists only of non-flagged nodes, and recall that our goal is to bound the length of P by a constant depending only on F and r . We refer to a maximal sequence of consecutive nodes of the same color along P as a *section* in this color (in the above argument, $\bar{Z}_1, \dots, \bar{Z}_b$ is a section in color s). Moreover, we call a section of P *internal* if it is neither the first nor the last section on P . By the argument above, for the last node \bar{Z} of an internal section and the last node Z of the preceding section on P , $x(\bar{Z})$ is higher on the walk \mathcal{W} than $x(Z)$. As the walk \mathcal{W} contains at most $r \cdot |\mathcal{S}(F)|$ different elements, the path P can have at most $r \cdot |\mathcal{S}(F)| - 1$ internal sections, and at most $r \cdot |\mathcal{S}(F)| + 1$ sections in total (including the first and last section). As each section consists of at most $v(F) + 1$ nodes, P consists of at most

$$(r \cdot |\mathcal{S}(F)| + 1)(v(F) + 1) =: p \quad (173)$$

nodes, proving that the length of P is indeed bounded by a constant depending only on F and r .

Since in total there are at most $v(F)/\varepsilon$ many flagged nodes in $\mathcal{T}(H')$, the depth of $\mathcal{T}(H')$ is bounded by

$$v(F)/\varepsilon + (v(F)/\varepsilon + 1)p =: t, \quad (174)$$

where the first term bounds the number of flagged nodes and the second term the number of non-flagged nodes. Consequently, we have

$$v(\mathcal{T}(H')) \leq 1 + r + r^2 + \dots + r^t \leq r^{t+1}.$$

Observing that every node of $\mathcal{T}(H')$ corresponds to at most $v(F)$ vertices of H' , we finally obtain that

$$v(H') \leq r^{t+1} \cdot v(F) \stackrel{(158), (173), (174)}{=} (v_{\max} - 1)/r.$$

This justifies (166) and concludes the proof. \square

6. PROOF OF THEOREM 4

We denote the board of the probabilistic process after i steps by G_i , where $0 \leq i \leq n$. We take the alternative view mentioned in Section 1.1, in which the random edges leading from a newly added vertex to previous vertices are generated at the moment this vertex is revealed instead of at the beginning of the process. (Recall that each edge is inserted with probability $p = p(n)$ independently from all other edges.) Thus G_i is an r -colored graph on i vertices, and the underlying uncolored graph is distributed as $G_{i,p}$.

Recall that all our asymptotic results are with respect to n , the number of vertices of G_n or $G_{n,p}$. We write $f(n) \ll g(n)$ if $f(n) = o(g(n))$, $f(n) \gg g(n)$ if $f(n) = \omega(g(n))$, and $f(n) \asymp g(n)$ if $f(n) = \Theta(g(n))$.

6.1. Lower bound. The crucial ingredient for the proof of the lower bound part of Theorem 4 is Lemma 35 from Section 5.

Proof of Theorem 4 (lower bound). Let $\theta^* = \theta^*(F, r)$ be defined as in Theorem 9, and let $\alpha^* = \alpha^*(F, r)$ be a sequence from the set $[r]^{r \cdot |\mathcal{S}(F)|}$ for which the minimum in the definition of $\Lambda_{\theta^*}(F, r)$

in (24) is attained. We show that the strategy $\text{PAINT}(F, r, \theta^*, \alpha^*)$ defined in Section 5.1 a.a.s. avoids F for all n steps of the process if

$$p \ll p_0(F, r, n) = n^{-1/m_1^*(F, r)} \stackrel{(16)}{=} n^{-\theta^*} . \quad (175)$$

By the choice of α^* and the definition in (24) we have that for all colors $s \in [r]$ and all vertex orderings $\pi \in \Pi(V(F))$ there is a subgraph $H \subseteq F$ such that

$$\lambda_{\theta^*}(H, w_{(H, \pi|_H, s)}) \leq \Lambda_{\theta^*}(F, r) \stackrel{(17)}{=} 0 . \quad (176)$$

According to Lemma 35 we then have for each such $(H, \pi|_H)$: if G_n contains a copy of $(H, \pi|_H)$ in color s , then it contains a graph K' with $v(K') \leq v_{\max}$ and $\mu_{\theta^*}(K') < 0$, or a graph H' with $v(H') \leq v_{\max}$ and

$$\mu_{\theta^*}(H') \stackrel{(111)}{\leq} \lambda_{\theta^*}(H, w_{(H, \pi|_H, s)}) \stackrel{(176)}{\leq} 0 .$$

This yields a family $\mathcal{W} = \mathcal{W}(F, \pi, s, r)$ of graphs W' satisfying $\mu_{\theta^*}(W') \leq 0$ and $v(W') \leq v_{\max}$ such that, deterministically, G_n contains a graph from \mathcal{W} if it contains a copy of (F, π) in color s . It follows that G_n contains a graph from

$$\mathcal{W}^* = \mathcal{W}^*(F, r) := \bigcup_{\substack{s \in [r] \\ \pi \in \Pi(V(F))}} \mathcal{W}(F, \pi, s, r)$$

if it contains a monochromatic copy of F .

Moreover, since no graph in \mathcal{W}^* has more than $v_{\max}(F, r, \theta^*(F, r), \alpha^*(F, r))$ vertices, the size of \mathcal{W}^* is bounded by a constant only depending on F and r . By the definition of $\mu_{\theta^*}()$ in (6) and the fact that $\mu_{\theta^*}(W') \leq 0$ for all $W' \in \mathcal{W}^*$, the expected number of copies of the (underlying uncolored) graphs from \mathcal{W}^* in $G_{n,p}$ is of order

$$\sum_{W' \in \mathcal{W}^*} n^{v(W')} p^{e(W')} \stackrel{(175)}{\ll} \sum_{W' \in \mathcal{W}^*} n^{\mu_{\theta^*}(W')} \leq |\mathcal{W}^*| \cdot n^0 = \Theta(1) .$$

It follows with Markov's inequality that a.a.s. $G_{n,p}$ contains no copy of any of the (underlying uncolored) graphs from \mathcal{W}^* . Consequently, a.a.s. G_n contains no copy of any of the graphs from \mathcal{W}^* and hence no monochromatic copy of F . This proves the claimed lower bound on the threshold of the probabilistic process.

To prove the second part of Theorem 4 it suffices to show that the strategy $\text{PAINT}(F, r, \theta^*, \alpha^*)$ is an optimal strategy for Painter in the deterministic two-player game, i.e., that it is a winning strategy in the game with density restriction d for any $d < m_1^*(F, r) = 1/\theta^*$ (we have already argued that this strategy can be implemented as a polynomial-time algorithm in Section 1.5 and Remark 34). Fix some $0 < d < 1/\theta^*$ and define $\theta := 1/d > \theta^*$. Suppose Painter plays according to the strategy $\text{PAINT}(F, r, \theta^*, \alpha^*)$ in the game with density restriction d and suppose for the sake of contradiction that the game ends with a monochromatic copy of F . Then as before it follows from Lemma 35 that the board contains a graph K' with

$$\mu_{\theta^*}(K') < 0 \quad (177)$$

or a graph H' with

$$\mu_{\theta^*}(H') \leq 0 \quad (178)$$

(note that H' contains at least one vertex and as a consequence of (178) and the definition in (6) also at least one edge; similarly, K' contains at least one edge as a consequence of (177)). Using that $\theta > \theta^*$ it follows from (177), (178) and the definition in (6) that in any case the board contains a graph W' (with $v(W') \geq 1$) satisfying $\mu_{\theta}(W') < 0$, or equivalently, $e(W')/v(W') > 1/\theta = d$, violating the given density restriction. \square

6.2. Upper bound. As in the proof of Lemma 8 we identify Builder's strategies in the deterministic two-player game with r colors with finite r -ary rooted trees, where each node at depth k of such a tree is an r -colored graph on k vertices, representing the board after the k -th step of the game.

Note that in this formalization, a given tree \mathcal{T} represents a generic strategy for Builder (in the deterministic game with r colors) that may or may not satisfy a given density restriction d , and that can be thought of as a strategy for the ' F -avoidance' game for any given graph F . We say that \mathcal{T} is a *winning strategy* for Builder in a specific F -avoidance game if and only if every leaf of \mathcal{T} contains a monochromatic copy of F . We say that a Builder strategy \mathcal{T} is a *legal strategy* in the game with density restriction d if and only if $e(H)/v(H) \leq d$ for every subgraph H of every node B in \mathcal{T} .

When we say that G_i , the board of the probabilistic process after i steps, contains a copy of some r -colored graph B (e.g. a node of some Builder strategy \mathcal{T}) we mean that there is a subgraph of G_i that is isomorphic to B as a *colored* graph.

The upper bound part of Theorem 4 is an immediate consequence of the following lemma.

Lemma 52 (Random process reproduces Builder strategy). *Let $r \geq 2$ be a fixed integer, let $d > 0$ be a fixed real number, and let \mathcal{T} represent an arbitrary legal strategy for Builder in the deterministic game with r colors and density restriction d .*

If $p \gg n^{-1/d}$, then regardless of the online coloring strategy employed, a.a.s. G_n contains a copy of a leaf of \mathcal{T} .

Proof of Theorem 4 (upper bound). By Theorem 3 there exists a legal winning strategy \mathcal{T} for Builder in the deterministic F -avoidance game with r colors and density restriction $d = m_1^*(F, r)$. As each leaf of \mathcal{T} contains a monochromatic copy of F , applying Lemma 52 to \mathcal{T} yields that if $p \gg p_0(F, r, n) = n^{-1/m_1^*(F, r)}$, then a.a.s. G_n contains a monochromatic copy of F , regardless of the online coloring strategy employed, which is exactly the upper bound statement of Theorem 4. \square

In order to prove Lemma 52, we shall show the following more technical statement by induction on k .

Lemma 53 (Random process reproduces Builder strategy step by step). *Let $r \geq 2$ be a fixed integer, let $d > 0$ be a fixed real number, and let \mathcal{T} represent an arbitrary legal strategy for Builder in the deterministic game with r colors and density restriction d .*

If $p \gg n^{-1/d}$, then for any integer $k \geq 1$ the following is true. Regardless of the online coloring strategy employed, a.a.s. one of the following two statements holds:

- G_n contains a copy of a leaf of \mathcal{T} , or
- there is a node B at depth k in \mathcal{T} such that G_n contains $\Omega(n^{v(B)} p^{e(B)})$ many copies of B .

The second property of Lemma 53 is meaningful since, due to the assumption that \mathcal{T} is a legal strategy for Builder in the game with density restriction d , we have

$$e(B)/v(B) \leq m(B) \leq d ,$$

which yields with $p \gg n^{-1/d} \geq n^{-v(B)/e(B)}$ that

$$n^{v(B)} p^{e(B)} \gg 1 .$$

Proof of Lemma 52. Set $k := \text{depth}(\mathcal{T}) + 1$ in Lemma 53. \square

It remains to prove Lemma 53.

Proof of Lemma 53. We proceed by induction on k . For the induction base $k = 1$, note that each of the r nodes B at depth 1 in \mathcal{T} consists simply of an isolated vertex, colored in one of the r available colors. Clearly, G_n contains at least $n/r = \Omega(n)$ copies of one of these by the pigeonhole principle.

For the induction step we employ a two-round approach. That is, we divide the process into two rounds of equal length $n/2$ (w.l.o.g. we assume n to be even) and analyze these two rounds separately. Denoting the vertices added throughout the process by v_1, \dots, v_n , the first round consists of adding the vertices $v_1, \dots, v_{n/2}$ together with the corresponding random edges. At the end of the first round, we thus obtain a graph $G_{n/2}$, to which we can apply the induction hypothesis and some standard random graph arguments. The second round consists of adding the vertices $v_{n/2+1}, \dots, v_n$ (together with the corresponding random edges). Using a variance calculation, we show that conditional on a ‘good’ first round, the second round turns out as claimed. (In fact, our argument does not make use of any edges added between vertices of the set $\{v_{n/2+1}, \dots, v_n\}$.)

By the induction hypothesis, if the graph $G_{n/2}$ does not contain a copy of a leaf of \mathcal{T} (in which case we are done), a.a.s. it contains a family of

$$M \asymp n^{v(B^\circ)} p^{e(B^\circ)} \quad (179)$$

copies of some graph B° corresponding to a non-leaf node at depth $k-1$ in \mathcal{T} . We label these copies B_i° , $1 \leq i \leq M$. Let B denote the graph obtained from B° by adding a new vertex v to it together with edges connecting v to B° as prescribed by Builder’s next move specified by \mathcal{T} (so v is uncolored in B , but assigning it one of the r available colors yields exactly one of the children of B° in \mathcal{T}).

For each copy B_i° , $1 \leq i \leq M$, and each vertex v_ℓ , $n/2 + 1 \leq \ell \leq n$, we fix a set $E_{i,\ell}$ of $\deg_B(v)$ many vertex pairs such that if the elements of $E_{i,\ell}$ are actual edges generated in the second round, then v_ℓ together with those edges completes B_i° to a copy of B . We let $Z_{i,\ell}$ be the indicator variable for the event that the elements of $E_{i,\ell}$ are generated as edges in the second round. Let

$$Z := \sum_{i=1}^M \sum_{\ell=n/2+1}^n Z_{i,\ell} ,$$

and note that by the pigeonhole principle at least Z/r many copies of one of the children of B° in \mathcal{T} are created. Thus the second condition of the lemma is satisfied if we show that a.a.s.

$$Z \asymp n^{v(B)} p^{e(B)} . \quad (180)$$

We will do so by the methods of first and second moment.

We clearly have

$$\Pr[Z_{i,\ell} = 1] = p^{|E_{i,\ell}|} = p^{\deg_B(v)} ,$$

and, conditioning on the first round satisfying the induction hypothesis,

$$\mathbb{E}[Z] = M \cdot n/2 \cdot p^{\deg_B(v)} \stackrel{(179)}{\asymp} n^{v(B)} p^{e(B)} . \quad (181)$$

In the following, we slightly abuse notation and write B also for the *uncolored* graph underlying B . Let \mathcal{D} denote the family of all (uncolored) graphs D that can be constructed by considering the union of two copies of B intersecting in at least two vertices, one of which must be the vertex v (we again slightly abuse notation in the following and refer to the corresponding vertex in each such graph D as v). For any $D \in \mathcal{D}$, we denote by D° the graph obtained by removing v from D .

To calculate the variance of Z , observe that the variables $Z_{i,\ell}$ and $Z_{j,\ell'}$ are independent whenever $\ell \neq \ell'$ or $B_i^\circ \cap B_j^\circ = \emptyset$. Hence such pairs can be omitted, and we have

$$\begin{aligned}
\text{Var}[Z] &= \sum_{i,j=1}^M \sum_{\ell,\ell'=n/2+1}^n (\mathbb{E}[Z_{i,\ell}Z_{j,\ell'}] - \mathbb{E}[Z_{i,\ell}]\mathbb{E}[Z_{j,\ell'}]) \\
&\leq \sum_{\substack{i,j=1,\dots,M: \\ B_i^\circ \cap B_j^\circ \neq \emptyset}} \sum_{\ell=n/2+1}^n \Pr[Z_{i,\ell} = 1 \wedge Z_{j,\ell} = 1] \\
&= \sum_{D \in \mathcal{D}} \sum_{\substack{i,j=1,\dots,M: \\ B_i^\circ \cap B_j^\circ = D^\circ}} \sum_{\ell=n/2+1}^n p^{|E_{i,\ell} \cup E_{j,\ell}|} \\
&\leq \sum_{D \in \mathcal{D}} M_{D^\circ} \cdot \Theta(1) \cdot np^{\deg_D(v)} ,
\end{aligned} \tag{182}$$

where M_{D° denotes the total number of copies of D° in (the underlying uncolored graph of) $G_{n/2}$. By definition of \mathcal{D} , each $D \in \mathcal{D}$ satisfies

$$\begin{aligned}
v(D^\circ) &= 2v(B) - v(J) - 1 , \\
e(D^\circ) &= 2e(B) - e(J) - \deg_D(v)
\end{aligned} \tag{183}$$

for some subgraph $J \subseteq B$. Moreover, since we assumed that \mathcal{T} is a legal strategy for Builder in the game with density restriction d , we have

$$e(J)/v(J) \leq m(B) \leq d ,$$

which yields with $p \gg n^{-1/d} \geq n^{-v(J)/e(J)}$ that

$$n^{v(J)} p^{e(J)} \gg 1 . \tag{184}$$

Thus the expected number of copies of D° in (the underlying uncolored graph of) $G_{n/2}$ is

$$\binom{n}{v(D^\circ)} \cdot \Theta(1) \cdot p^{e(D^\circ)} \stackrel{(183)}{\asymp} n^{2v(B)-v(J)-1} p^{2e(B)-e(J)-\deg_D(v)} \stackrel{(184)}{\ll} n^{2v(B)-1} p^{2e(B)-\deg_D(v)} .$$

and Markov's inequality implies that

$$M_{D^\circ} \ll n^{2v(B)-1} p^{2e(B)-\deg_D(v)} \tag{185}$$

a.a.s. As moreover the number of graphs in \mathcal{D} is bounded by a constant depending only on \mathcal{T} , a.a.s. (185) holds for all $D \in \mathcal{D}$ simultaneously.

Thus, conditioning on the first round satisfying the induction hypothesis (cf. (179)), and (185) for all $D \in \mathcal{D}$, we obtain from (182) that

$$\text{Var}[Z] \stackrel{(185)}{\ll} \sum_{D \in \mathcal{D}} \left(n^{v(B)} p^{e(B)} \right)^2 \stackrel{(181)}{\asymp} \mathbb{E}[Z]^2 .$$

Chebyshev's inequality now yields that a.a.s. the second round satisfies (180). This implies that there is at least the claimed number of copies of one of the children of B° in G_n , as discussed. \square

ACKNOWLEDGEMENT

We thank Michael Belfrage for many inspiring discussions about online Ramsey games.

REFERENCES

- [1] currently <http://www.as.inf.ethz.ch/muetze>, <http://www.as.inf.ethz.ch/rast>, and <http://www.mpi-inf.mpg.de/~rspoehel>.
- [2] D. Achlioptas and E. Friedgut. A sharp threshold for k -colorability. *Random Structures Algorithms*, 14(1):63–70, 1999.
- [3] N. Alon and M. Krivelevich. The concentration of the chromatic number of random graphs. *Combinatorica*, 17(3):303–313, 1997.
- [4] N. Alon, M. Krivelevich, and B. Sudakov. List coloring of random and pseudo-random graphs. *Combinatorica*, 19(4):453–472, 1999.
- [5] J. Balogh and J. Butterfield. Online Ramsey games for triangles in random graphs. *Discrete Mathematics*, 310(24):3653–3657, 2010.
- [6] J. Beck. There is no fast method for finding monochromatic complete subgraphs. *J. Combin. Theory Ser. B*, 34(1):58–64, 1983.
- [7] J. Beck. Achievement games and the probabilistic method. In *Combinatorics, Paul Erdős is eighty, Vol. 1*, Bolyai Soc. Math. Stud., pages 51–78. János Bolyai Math. Soc., Budapest, 1993.
- [8] M. Belfrage, T. Mütze, and R. Spöhel. Probabilistic one-player Ramsey games via deterministic two-player games. Submitted.
- [9] T. Bohman, A. Frieze, and B. Sudakov. The game chromatic number of random graphs. *Random Structures Algorithms*, 32(2):223–235, 2008.
- [10] B. Bollobás. Threshold functions for small subgraphs. *Math. Proc. Cambridge Philos. Soc.*, 90(2):197–206, 1981.
- [11] B. Bollobás. The chromatic number of random graphs. *Combinatorica*, 8(1):49–55, 1988.
- [12] B. Bollobás and A. Thomason. Generalized chromatic numbers of random graphs. In *Proceedings of the Sixth International Seminar on Random Graphs and Probabilistic Methods in Combinatorics and Computer Science, “Random Graphs ’93” (Poznań, 1993)*, volume 6, pages 353–356, 1995.
- [13] J. Butterfield, T. Grauman, W. Kinnersley, K. Milans, C. Stocker, and D. West. On-line Ramsey theory for bounded degree graphs. Submitted, 2011.
- [14] A. Coja-Oghlan, K. Panagiotou, and A. Steger. On the chromatic number of random graphs. *J. Combin. Theory Ser. B*, 98(5):980–993, 2008.
- [15] D. Conlon. On-line Ramsey numbers. *SIAM J. Discrete Math.*, 23(4):1954–1963, 2009/10.
- [16] E. Friedgut. Hunting for sharp thresholds. *Random Structures Algorithms*, 26(1-2):37–51, 2005.
- [17] E. Friedgut, Y. Kohayakawa, V. Rödl, A. Ruciński, and P. Tetali. Ramsey games against a one-armed bandit. *Combin. Probab. Comput.*, 12(5-6):515–545, 2003. Special issue on Ramsey theory.
- [18] E. Friedgut and M. Krivelevich. Sharp thresholds for certain Ramsey properties of random graphs. *Random Structures Algorithms*, 17(1):1–19, 2000.
- [19] G. R. Grimmett and C. McDiarmid. On colouring random graphs. *Math. Proc. Cambridge Philos. Soc.*, 77:313–324, 1975.
- [20] J. A. Grytczuk, M. Hałuszczak, and H. A. Kierstead. On-line Ramsey theory. *Electron. J. Combin.*, 11(1):Research Paper 60, 10 pp. (electronic), 2004. Paper number later changed by the publisher from 60 to 57.
- [21] J. A. Grytczuk, H. A. Kierstead, and P. Prałat. On-line Ramsey numbers for paths and stars. *Discrete Math. Theor. Comput. Sci.*, 10(3):63–74, 2008.
- [22] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
- [23] R. Kang and C. McDiarmid. The t -improper chromatic number of random graphs. *Combin. Probab. Comput.*, 19(1):87–98, 2010.
- [24] H. A. Kierstead and G. Konjevod. Coloring number and on-line Ramsey theory for graphs and hypergraphs. *Combinatorica*, 29(1):49–64, 2009.
- [25] M. Krivelevich and B. Patkós. Equitable coloring of random graphs. *Random Structures Algorithms*, 35(1):83–99, 2009.
- [26] A. Kurek and A. Ruciński. Globally sparse vertex-Ramsey graphs. *J. Graph Theory*, 18(1):73–81, 1994.
- [27] A. Kurek and A. Ruciński. Two variants of the size Ramsey number. *Discuss. Math. Graph Theory*, 25(1-2):141–149, 2005.
- [28] P.-S. Loh and B. Sudakov. On the strong chromatic number of random graphs. *Combin. Probab. Comput.*, 17(2):271–286, 2008.
- [29] T. Łuczak. The chromatic number of random graphs. *Combinatorica*, 11(1):45–54, 1991.
- [30] T. Łuczak. A note on the sharp concentration of the chromatic number of random graphs. *Combinatorica*, 11(3):295–297, 1991.
- [31] T. Łuczak, A. Ruciński, and B. Voigt. Ramsey properties of random graphs. *J. Combin. Theory Ser. B*, 56(1):55–68, 1992.

- [32] M. Marciniszyn and R. Spöhel. Online vertex-coloring games in random graphs. *Combinatorica*, 30(1):105–123, 2010. An extended abstract appeared in the proceedings of SODA '07.
- [33] M. Marciniszyn, R. Spöhel, and A. Steger. Online Ramsey games in random graphs. *Combin. Probab. Comput.*, 18(Special Issue 1-2):271–300, 2009.
- [34] M. Marciniszyn, R. Spöhel, and A. Steger. Upper bounds for online Ramsey games in random graphs. *Combin. Probab. Comput.*, 18(Special Issue 1-2):259–270, 2009.
- [35] D. W. Matula. Expose-and-merge exploration and the chromatic number of a random graph. *Combinatorica*, 7(3):275–284, 1987.
- [36] C. McDiarmid. On the chromatic number of random graphs. *Random Structures Algorithms*, 1(4):435–442, 1990.
- [37] T. Mütze and R. Spöhel. On the path-avoidance vertex-coloring game. Submitted.
- [38] P. Prałat. A note on small on-line Ramsey numbers for paths and their generalization. *Australas. J. Combin.*, 40:27–36, 2008.
- [39] P. Prałat. $\overline{\mathcal{R}}(3, 4) = 17$. *Electron. J. Combin.*, 15(1):Research Paper 67, 13, 2008.
- [40] E. Scheinerman. Generalized chromatic numbers of random graphs. *SIAM J. Discrete Math.*, 5(1):74–80, 1992.